

ADTSC Science Highlights 2010

Associate Directorate for Theory, Simulation, and Computation

Preface

The Theory, Simulation, and Computation (TSC) Directorate is central to the huge national need for new generations of ideas, concepts, and methodologies to improve the fidelity, reliability, certainty and usability of tools to guide and interpret experiments, and provide prediction and control for complex phenomena and systems: the integration of information, science, and technology between data and prediction. Working with teams across the Laboratory, TSC undertakes major interdisciplinary challenges of integrating theory, modeling, simulation, high-performance computing, and visualization with experimental and other data to expand LANL's predictive and uncertainty quantification capability. The core nuclear weapons program depends critically on the viability of this approach. However, virtually every current and emerging mission at LANL also relies centrally on a similar integrated capability: energy and global security, manufacturing science, and many discipline frontiers from biology to materials science to cosmology.

Each year, TSC staff select a collection of brief topical reports on their recent research results. This annual collection is not intended to be complete. Rather it aims to demonstrate, through these timely examples, the interdisciplinary energy and progress across many scientific and programmatic frontiers. The collection chosen this year represents work for the National Nuclear Security Administration and also presents examples from our “work with others” portfolio. This collection is representative of our extensive outreach activities to universities, industry, and other research laboratories, and is a basis for providing the best science and scientists to both current and future national security missions at LANL.

Alan R. Bishop

Associate Director, Theory, Simulation, and Computation Directorate

Alan R. Bishop

Associate Director

Paul J. Dotson

Deputy Associate Director

Andrew B. White

*Deputy Associate Director for
High-Performance
Computing*

Audrey L. Archuleta

ADTSC Chief of Staff

Stephen R. Lee

CCS Division Leader

John F. Morrison

HPC Division Leader

Antonio Redondo

T Division Leader

Publications Team

Sharon Mikkelsen

Design and Coordination

Susanne King,

Kathy Pallis

Editing

Guadalupe D. Archuleta

Print Coordinator

ADTSC web site: <http://www.lanl.gov/orgs/adtsc>





2010 ADTSC Science Highlights Table of Contents

Applied Mathematics and Fluid Dynamics

Probability Density Function Method for Variable Density Turbulence	2
Jozsef Bakosi, XCP-1, J. R. Ristorcelli, XCP-4	
Computational Modeling of High Data-rate Laser Communication in Turbulent Atmosphere	4
Gennady P. Berman, T-4/T-CNLS; Alan R. Bishop, ADTSC; Boris M. Chernobrod, T-4; O.O. Chumak, T-4, National Academy of Sciences of Ukraine, Institute of Physics; Vyacheslav N. Gorshkov, T-4/T-CNLS and National Technical University of Ukraine, and National Academy of Sciences of Ukraine, Institute of Physics; S.V. Torous, National Technical University of Ukraine	
Adaptive Finite Element Methods for Turbulent and Reactive Flow	6
David Carrington, David Torres, T-3	
Development of a New Adaptive Mesh Refinement Framework for Multiphysics Simulations	8
William W. Dai, HPC-5	
VPIC–Magnetic Reconnection	10
William S. Daughton, XCP-6 , Vadim Roytershteyn, T-5; Lin Yin, Brian J. Albright, XCP-6; Ben Bergen, CCS-2; Kevin J. Bowers, XCP-6	
Accurate Estimation of Geometrical Properties of a Volume-of-Fluid Interface on Nonuniform Rectangular Grids	12
Marianne M. Francois, CCS-2; Blair K. Swartz, T-5	
Conformal Refinement and Coarsening of Unstructured Quadrilateral Meshes	14
Rao V. Garimella, T-5	
Effective Shear Viscosity of Active Suspensions at Moderate Concentrations	16
Vitaliy Gyrya, Pennsylvania State University; Konstantin Lipnikov, T-5; Igor Aronson, Argonne National Laboratory; Leonid Berlyand, Pennsylvania State University	
High Atwood Number Rayleigh-Taylor Instability	18
Daniel Livescu, Mark R. Petersen, CCS-2; Steven L. Martin, Patrick S. McCormick, CCS-1	



Direct Numerical Simulations of Reacting Compressible Turbulence	20
Daniel Livescu, Jamaludin Mohd-Yusof, Timothy M. Kelley, CCS-2	
Shock Wave Structure for an Ionized Plasma	22
Thomas O. Masser, John G. Wohlbiel, Robert B. Lowrie, CCS-2	
A Realistic and Stable Method of Forcing for Compressible Isotropic Turbulence.....	24
Mark R. Petersen, Daniel Livescu, Sumner Dean, Jamaludin Mohd-Yusof, CCS-2	
On the Early-time Kinetic Energy Profile for the Rayleigh-Taylor Instability	26
Bertrand Rollin, Malcolm J. Andrews, CCS-2	
Improved Interface Treatment in Arbitrary Lagrangian-Eulerian Simulations using the Moment-of-Fluid Method	28
M. Kucharik, FNSPE, Czech Technical University in Prague; Rao V. Garimella, Samuel P. Schofield, Mikhail J. Shashkov, T-5	
Turbulent Mixing Simulations with Verification and Validation	30
James Glimm, David H. Sharp, ADTSC	
High-order Divergence-free MHD Solver for Turbulence Simulation	32
Shengtai Li, T-5	
A New Monotone Finite Volume Method for Advection-Diffusion Problems	34
Konstantin Lipnikov, Daniil Svyatskiy, T-5; Yuri Vassilevki, Institute of Numerical Mathematics of the Russian Academy of Sciences, Russia	
Understanding the Nonlinear Physics of Laser-Plasma Interaction through "At-scale" Plasma Kinetic Simulations on Roadrunner	36
Lin Yin, Brian J. Albright, XCP-6; Kevin J. Bowers, XCP-6; Ben Bergen, CCS-2	
Robust Mesh Motion Based on Monge-Kantorovich Equidistribution.....	38
Luis Chacón, Oak Ridge National Laboratory; Gian Luca Delzanno, John M. Finn, T-5	



A Particle-in-Cell Code with Arbitrary Curvilinear Mesh	40
Chris A. Ficht, University of New Mexico; John M. Finn, T-5	
Robust Multigrid with Cell-based Coarsening by a Factor of Three	42
Joel E. Dendy, David J. Moulton, T-5	
Numerical Simulation of Single Droplet Dynamics with Species Diffusion for Fuel Reprocessing	44
Marianne M. Francois, Neil N. Carlson, CCS-2	

Astrophysics and Cosmology

Hybrid Petacomputing Meets Cosmology: The Roadrunner Universe Project.....	48
Salman Habib, T-2; Adrian Pope, ISR-1/T-2/CCS-6; Zarija Lukić, T-2; David Daniel, CCS-1; Patricia Fasel, CCS-3; Katrin Heitmann, ISR-1; Nehal Desai, CCS-1 and Aerospace; Chung-Hsing Hsu, CCS-1 and Oak Ridge National Laboratory; Lee Ankeny, HPC-1; Graham Mark, CCS-3; Suman Bhattacharya, T-2; James Ahrens; CCS-1	
Cosmic Emulation: The Universe as a <i>Black Box</i>	50
Earl Lawrence, CCS-6; Katrin Heitmann, ISR-1; David Higdon, CCS-6; Christian Wagner, ISR-1; Martin White, University of California, Berkeley; Salman Habib, T-2; Brian Williams, CCS-6	
Modeling Progenitors and Spectra from White Dwarf Collisions: Thermonuclear and Core-Collapse Supernovae	52
Chris L. Fryer, CCS-2; Aimee Hungerford, XTD-6; Gabriel Rockefeller, CCS-2; Bob Greene, XCP-2; Frank Timmes, Cody Raskin, Arizona State University	
Molecular Dynamics Simulation of Multispecies Neutron Star Crust Conductivity	54
Sanjib Gupta, Jerome Daligault, T-5	
The Roadrunner Universe Project: Baryon Acoustic Oscillations in the Intergalactic Medium	56
Adrian Pope, ISR-1/T-2/CCS-6; Katrin Heitmann, ISR-1; Salman Habib, Zarija Lukić, T-2; Patricia Fasel, CCS-3; David Daniel, CCS-1; Martin White, Jordan Carlson, University of California, Berkeley	
Fast and Accurate Simulations of Proto-planet Migration in Disks	58
Shengtai Li, T-5; Hui Li, T-2	



Using Roadrunner to Model Shock Breakout in Supernovae and National Ignition Facility Experiments	60
Timothy M. Kelley, Todd Urbatsch, CCS-2; Aimee Hungerford, XTD-6; Gabriel Rockefeller, Chris L. Fryer, Paul J. Henning, Jeff Densmore, CCS-2; Barbara DeVolder, XCP-6	

Origin of Magnetic Fields in Galaxy Clusters	62
Hao Xu, Hui Li, T-2	

Atomic, Nuclear, and High-Energy Physics

Ab initio Parallel Replica Molecular Dynamics Method	66
Elena Jakubikova, Danny Perez, Artur F. Voter, Richard L. Martin, Enrique R. Batista, T-1	

Simulation of X-ray Diffraction Patterns	68
John L. Barber, T-1	

Linear Scaling Quantum Molecular Dynamics for Organic Molecules	70
Edward Sanville, Nicolas Bock, Matt Challacombe, Anders M. N. Niklasson, Marc Cawkwell, T-1	

Deep Interference Minima in Differential Cross Sections for Electron-impact Ionization	72
James Colgan, T-1	

Chemistry and Biology

HIV vs Supercomputing	76
Tanmoy Bhattacharya, T-2; Marcus Daniels, S. Gnanakaran, Bette Korber, T-6	

Conformational Variability of Soluble Cellulose Oligomers	78
S. Gnanakaran, T-6; Tongye Shen, University of Tennessee and ORNL; Paul Langan, B-8; Alfred D. French, Glenn Johnson, Southern Regional Research Center, USDA	

Probing the Dark Matter of the Genome: Mechanistic Studies of Noncoding RNAs	80
Karissa Y. Sanbonmatsu, Scott P. Hennelly, T-6	



Building Models of Protein Complexes Using the Known Protein-protein Interfaces: Structural Model of PhoB Dimeric Complex in its Active State	82
Chang-Shung Tung, T-6	

Quantitative Insights into Gene Regulation	84
Robert G. Martin, Judah L. Rosner, National Institutes of Health; Michael E. Wall, CCS-3	

Climate, Atmospheric, and Earth Systems Modeling

Observing and Predicting Changes in Greenland's Surface Mass Balance	88
Sebastian H. Mernild, CCS-2; William H. Lipscomb, T-3	

Tracer Advection for Ocean and Atmospheric Flows	90
Robert B. Lowrie, CCS-2; Todd D. Ringler, T-3	

Subbasin-based Parallel Hydrological Models for Semi-arid Regions	92
Sue Mniszewski, Patricia Fasel, CCS-3; Everett Springer, STBPO-PRM; Enrique Vivoni, Arizona State University; Amanda White, New Mexico Institute of Mining and Technology	

Bridging Scales in the Earth's Climate System through Multiresolution Modeling	94
Todd D. Ringler, T-3	

Information Science and Technology

PLFS: A Checkpoint Filesystem for Parallel Applications	98
John Bent, HPC-5; Garth Gibson, Carnegie Mellon University; Gary Grider, HPC-DO; Ben McClelland, HPC-3; Paul Nowoczynski, Pittsburgh Supercomputing Center; James Nunez, Milo Polte, Meghan Wingate, HPC-5	

Data Mining in Radiation Portal Monitoring.....	100
Tom Burr, Michael Hamada, CCS-6; Nicolas Hengartner, CCS-2; Kary Myers, Richard Picard, CCS-6	



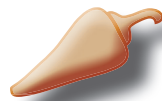
Fast Algorithms for Image Reconstruction from Very Few Data	102
Rick Chartrand, T-5	
COTS Parallel Archive Integration Experiences	104
Gary Grider, HPC-DO; Hsing-Bung Chen, HPC-5; Cody Scott, Milton Turley, Aaron Torres, Kathryn Sanchez, John Bremer, HPC-3	
Report on Current Issues in Resilience	106
Nathan DeBardeleben, HPC-5	
Statistical Models to Support Assessment and Decision-making in Stockpile Stewardship	108
Aparna V. Huzurbazar, Christine Anderson-Cook, Elizabeth Kelly, Michael Fugate, Todd Graves, Michael Hamada, Geraldyn Hemphill, David Higdon, Richard Klamann, Lisa Moore, Richard Picard, Lawrence Ticknor, Joanne Wendelberger, Scott Vander Wiel, Brian Williams, CCS-6	
Using the MILTS-Cyber Tool to Analyze Large-scale Cyber Threats	110
Guanhua Yan, Stephan Eidenbenz, CCS-3	

Materials Science

Systematic Study of Modifications to Ru(II)-polypyridine Dyads for Electron Injection Enhancement.....	114
Elena Jakubikova, Richard L. Martin, Enrique R. Batista, T-1; Robert C. Snoeberger III, Victor S. Batista, Yale University	
Modeling and Characterization of Grain Scale Strain Distribution in Polycrystalline Tantalum.....	116
Curt A. Bronkhorst, T-3; Amy R. Ross, MST-16; Benjamin L. Hansen, Hashem M. Mourad, T-3; Ellen K. Cerreta, John F. Bingert, MST-8	
Nonequilibrium Volumetric Response of Shocked Polymers	118
Bradford E. Clements, T-1	
Material Dynamics at Extreme Conditions	120
Timothy C. Germann, T-1; Sheng-Nian Luo, P-24; Sriram Swaminarayan, CCS-2	



Modeling Metallic Single Crystal Plastic Hardening through the Evolution of Dislocation Subgrain Structure	122
Benjamin L. Hansen, Curt A. Bronkhorst, T-3; M. Ortiz, California Institute of Technology	
Statistical Modeling for U-Nb Aging and Lifetime Prediction	124
Geraldyn Hemphill, CCS-6; Robert Hackenberg, MST-6	
A Computational Study of the Ballistic Performance of Composite Materials	126
Balaji Jayaraman, Xia Ma, Paul T. Giguere, Duan Z. Zhang, T-3	
Extended Lagrangian Born-Oppenheimer Molecular Dynamics	128
Anders M. N. Niklasson, T-1; Peter Steneteg, Linköping University, Sweden; Marc Cawkwell, T-1	
Simulating the Mechanical Behavior of Metallic Nanowires over Experimentally Accessible Timescales on Roadrunner.....	130
Danny Perez, T-1; Chun-Wei Pao, Academia Sinica, Taiwan; Sriram Swaminarayan, CCS-2; Arthur F. Voter, T-1	



Funding Acknowledgment Cross Reference

Department of Energy (DOE)

DOE, Office of Science	40, 52
Advanced Scientific Computing Research (ASCR)	16, 34, 42
Fusion Energy Sciences	62
High Energy Physics	50
Office of Basic Energy Sciences (BES)	130
Regional and Global Climate Modeling Program	90, 94
Scientific Discovery through Advanced Computing Program (SciDAC)	88, 98

DOE, Energy Efficiency and Renewable Energy (EERE)

Advanced Combustion Engine Research and Development	6
---	---

DOE, Office of Biological and Environmental Research (OBER)

Climate Change Prediction Program	88
---	----

DOE, Office of Nuclear Energy

Global Nuclear Energy Partnership	132
Nuclear Energy Advanced Modeling and Simulation (NEAMS) Program	44
Science and Technology, Advanced Fuel Cycle Initiative (AFCI)	132

National Nuclear Security Administration (NNSA)

DOE, NNSA Defense Programs and Core Surveillance	108
DOE, NNSA Enhanced Surveillance Campaign	108, 124
DOE, NNSA, Predictive Science Academic Alliance Program (PSAAP)	30
DOE, NNSA Advanced Simulation and Computing (ASC) Program	2, 8, 10, 22, 24, 70, 72, 88, 116, 122, 124

Department of Homeland Security (DHS).....

DHS, National Infrastructure Simulation and Analysis Center (NISAC)	110
---	-----

Department of Defense (DoD)

Army Research Office (ARO)	30
DoD and DOE Joint Munitions Technology Development Program	108, 118, 126



Defense Advanced Research Projects Agency (DARPA)	
Multi-aperture Geosynchronous Imaging Project	102
LANL Directed Research and Development Program (LDRD)	8, 10, 12, 14, 18, 20, 26, 28, 33,
.....	36, 38, 48, 50, 52, 54, 56, 58, 62, 66,
.....	70, 76, 78, 80, 82, 84, 92, 102, 114, 120, 128, 130
LANL Institute of Geophysics and Planetary Physics (IGPP)	62, 88
LANL Institutional Computing Program	62, 72, 104
LANL MaRIE Capability Development	68
National Institutes of Health (NIH)	
NIH American Recovery and Reinvestment Act (ARRA).....	80
NIH Intramural Research Program	84
National Science Foundation (NSF)	66, 98
NSF Theoretical Reseach.....	52
Office of Naval Research (ONR)	4
National Aeronautics and Space Administration (NASA)	
NASA Theoretical Astrophysics Program.....	50, 56



Organizational Abbreviations (for this publication only)

ADTSC

Associate Directorate Theory, Simulation, and Computation

Bioscience Division

B-8 Bioenergy and Environmental Science

Computer, Computational, and Statistical Sciences Division

CCS-1 Computer Science for HPC

CCS-2 Computational Physics and Methods

CCS-3 Information Sciences

CCS-6 Statistical Sciences

High Performance Computing Division

HPC-1 Scientific Software Engineering

HPC-3 High Performance Computer Systems

HPC-5 System Integration

HPC-DO High Performance Computing Division Office

International, Space, and Response Division

ISR-1 Space Science and Applications

Materials Science and Technology Division

MST-6 Materials Technology - Metallurgy

MST-8 Structure/Property Relations

MST-16 Nuclear Materials Science

Physics Division

P-24 Plasma Physics

Science and Technology Base Program Office

STBPO-PRM Peer Review and Metrics

Theoretical Division

T-1 Physics and Chemistry of Materials

T-2 Nuclear and Particle Physics, Astrophysics and Cosmology

T-3 Fluid Dynamics and Solid Mechanics

T-4 Physics of Condensed Matter & Complex Systems

T-5 Applied Mathematics and Plasma Physics

T-6 Theoretical Biology and Biophysics

T-CNLS Center for Nonlinear Studies

Applied Computational Physics Division

XCP-1 Lagrangian Codes

XCP-2 Eulerian Codes

XCP-4 Methods and Algorithms

XCP-6 Computational Physics

Applied Theoretical Design Division

XTD-6 Theoretical Design

ANL Argonne National Laboratory

NIH National Institutes of Health

LLNL Lawrence Livermore National Laboratory

ORNL Oak Ridge National Laboratory

SRRC Southern Regional Research Center, USDA

USDA U.S. Department of Agriculture



Applied Mathematics and Fluid Dynamics

LANL has a rich history in the development of novel applied mathematics and computational fluid dynamics reaching back to the earliest days of the Laboratory. This history has always had a special emphasis on modeling and understanding challenging physical situations including turbulent flow, shock physics, magnetohydrodynamics, hydrodynamic instabilities, and plasma physics. The tradition has continued and is described here with several articles discussing recent work on these topics. There are brief articles on turbulence, numerical methods, Rayleigh-Taylor instability, and plasma simulations, and verification and validation.

Applied Mathematics and Fluid Dynamics

Probability Density Function Method for Variable Density Turbulence

Jozsef Bakosi, XCP-1; J.R. Ristorcelli, XCP-4

Variable density (VD) turbulence, as a result of the turbulent mixing of fluids with different densities, differs from incompressible shear-driven turbulence in some very fundamental ways. VD flows are present in many geophysical, astrophysical, and engineering situations. These flows are typically driven by large pressure gradients (e.g., gravity), in which the different inertias of the two different-density fluids results in large differential fluid accelerations. If the fluid configuration is unstable, i.e., the body force is directed opposite to the density gradient, small perturbations of the initial interface between the fluids grow and eventually give rise to a highly turbulent flow. This phenomenon is known as the Rayleigh-Taylor (RT) instability and its understanding and predictability are vitally important in the simulation of unstably stratified oceanic and atmospheric flows, as well as inertial confinement fusion and supernova explosions.

The computational tools used to calculate these flows are statistical turbulence closures and direct numerical simulation (DNS). DNS is an extremely useful research tool: it resolves all dynamical scales of fluid motion and provides abundant information on the flow evolution. However, its use for practical purposes is limited since it requires substantial computing power. Therefore, there is a need for statistical models that capture only certain important features of these flows, at a computational price of orders of magnitude smaller than that of DNS.

Compared with moment closures of turbulence, which traditionally solve for the first (or the first two) statistical moment of the fluctuating flow variables, probability density function (PDF) methods [1] provide a higher level statistical description by solving for the PDF, containing information on all moments. Modeling assumptions are still required, but only for two-point quantities,

such as energy dissipation and molecular diffusion.

A PDF method for turbulent hydrodynamics and active scalar mixing [2,3] has been developed for VD turbulence that has the following features:

- Provides the full time-accurate evolution of the joint PDF of the fluid density and velocity in an RT-unstable flow, starting from a quiescent state and transitioning to highly nonequilibrium turbulence
- Captures the essential features of VD turbulence and mixing, such as the mixing asymmetry due to the possibly vastly different inertia of the fluids
- Correctly represents the anisotropy of the Reynolds stresses, important in predicting directional effects of turbulent mixing and the production of kinetic energy

To date statistical models cannot predict any of the above features. In addition, compared with most turbulence models developed for equilibrium flows, this method [3] surprisingly captures the time-accurate representation of the highly nonlinear process of transition to turbulence.

Compared with moment closures, PDF methods require no explicit modeling for the important physical processes of advection, chemical reactions, and the effects of mass flux: these terms appear in closed form in the sample space. The correct representation of these effects are crucial in predicting turbulent flows in general, and VD flows in particular.

In VD flows, where the densities of the mixing fluids are vastly different, the effects of large density fluctuations cannot be neglected. This is unlike the much studied and simpler class of mixing that occurs in the Boussinesq fluid case in which the densities are commensurate. In VD flows the large density differences give rise to cubic nonlinearities in the Navier-Stokes equation and mathematically very different molecular mixing rates, resulting in several non-Boussinesq effects. Examples are the mixing asymmetry, indicated by the nonzero skewness of the density PDF and the dynamic evolution of the mean pressure gradient [4].

Compared with the Boussinesq case, these features pose significant

challenges for modeling. Our model is the first statistical turbulence closure to capture these recently discovered [4] VD effects.

The numerical solution of the PDF transport equation amounts to following a large number of Lagrangian particles in a Monte Carlo fashion, whose properties are governed by an equivalent system of stochastic differential equations. Thus the software implementation of the method is efficient and highly scalable.

As an example of the high-level statistical information that can be extracted from the joint PDF computed by the method, the statistical density distribution, as it evolves in time at the center of the RT mixing layer, is plotted in Fig. 1. Six different points in time are displayed for a case when the density ratio of the mixing fluids is 1:3. For validation, the density PDFs that are extracted from DNS [4] are also plotted at the same times. Figure 1 shows the excellent prediction of the full PDF at all times, representing the correct mixing state and mixing asymmetry.

The knowledge of the full PDF, provided by the method, is invaluable in the correct prediction of coupled turbulence-radiation interactions in inertial confinement fusion or astrophysical applications, combustion calculations with detailed chemistry or uncertainty quantification of hazardous releases in the atmospheric boundary layer.

For more information contact Jozsef Bakosi at jbakosi@lanl.gov.

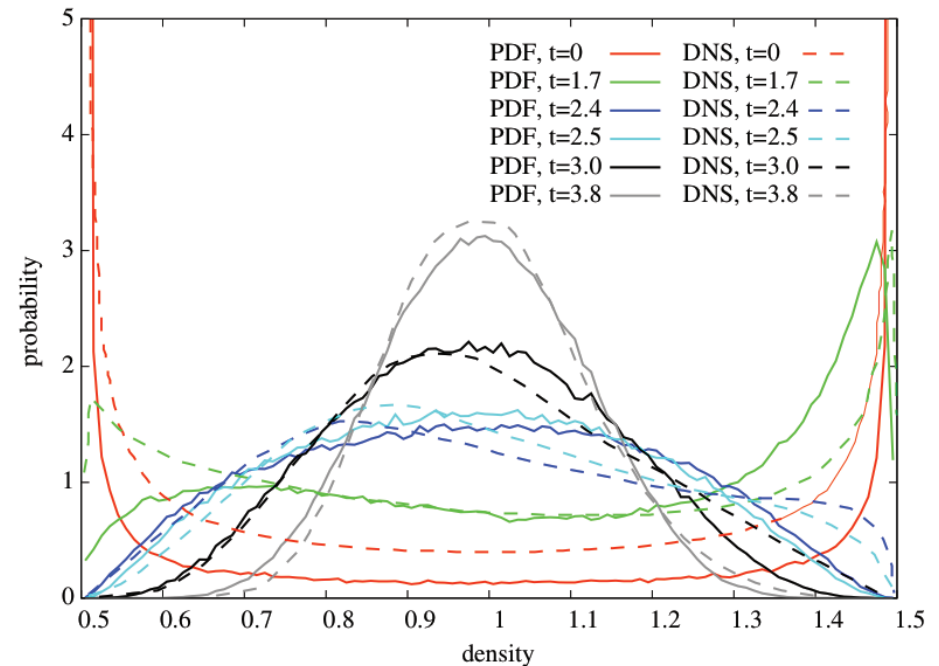


Fig. 1. Time evolution of the probability density function (PDF) of the fluid density at the center of a Rayleigh-Taylor mixing layer at the initial density ratio 1:3. Solid lines – PDF model calculation, dashed lines – direct numerical simulations (DNS) [4]. The flow starts from a quiescent state at $t=0$, corresponding to an approximate double-delta distribution, representing unmixed fluids. Molecular mixing, as a consequence of turbulent stirring, quickly diffuses the PDF. Variable-density (non-Boussinesq) effects already appear at this density ratio due to large differential accelerations. The higher density fluid mixes slower than the lower density one. This results in a skewed distribution, indicated by the green line at $t=1.7$.

- [1] W. Kollmann, *Theoret. Comput. Fluid Dyn.* **1**, 249 (1990).
- [2] J. Bakosi, J.R. Ristorcelli, *Transported Beta-PDF Model For Turbulent Mixing* (submitted to *Phys. Fluids* 2010).
- [3] J. Bakosi, J.R. Ristorcelli, *A probability Density Function Model for Variable-density Pressure-gradient Driven Turbulence and Mixing* (submitted to *Phys. Fluids* 2010).
- [4] D. Livescu, J.R. Ristorcelli *J. Fluid Mech.* **605**, 145 (2008).

Funding Acknowledgments

DOE, NNSA Advanced Simulation and Computing (ASC) Program

Computational Modeling of High Data-rate Laser Communication in Turbulent Atmosphere

Gennady P. Berman, T-4/T-CNLS; Alan R. Bishop, ADTSC; Boris M. Chernobrod, T-4; O.O. Chumak, T-4, National Academy of Sciences of Ukraine, Institute of Physics; Vyacheslav N. Gorshkov, T-4/T-CNLS, National Technical University of Ukraine, and National Academy of Sciences of Ukraine, Institute of Physics; S.V. Torous, National Technical University of Ukraine

The scintillations of intensity (SI) of a laser beam caused by the influence of atmospheric turbulence are a major obstacle for gigabit data rates and long-distance optical communications [1,2]. In the process of propagation of the laser beam through turbulent atmosphere with a given strength, C_n^2 , the signal (intensity, I^l) at the detector depends on the state (realization) of atmosphere, l .

One of the main factors contributing to the

$$SI, \sigma^2 = \langle (I^l)^2 \rangle_l / \langle I^l \rangle_l^2 - 1$$

is a beam fragmentation, which results in a decay of the initial laser beam into a set of separated beams. Some of these beams may not be detected if the size of detector is smaller than the characteristic distance between the beams. Moreover, depending on l , the coherent beam initially oriented along the z-axis deviates as a whole in the x-y plane. This deviation is characterized by the beam wandering [3], (Fig. 1).

$$\vec{r}_w^l(z) = \int \vec{r} I^l(x, y, z) dx dy / \int I^l(x, y, z) dx dy.$$

The integral dependence of the intensity, I^l , on r_w^l contributes to the high level of scintillations, σ^2 , at the detector.

It is known that a partially coherent beam (PCB) in combination with a slow-time-response detector leads to a significant reduction of the SI. However, this time-averaging (slow detector) technique is incompatible with gigabit data rate. Recently, we proposed a new method for the reduction of scintillations that combines a time averaging of a PCB with a spectral encoding [4-7]. In our approach, the information is encoded in the form of dips in the spectrum of subpicosecond coherent laser pulses.

The key component of the proposed technique is a phase diffuser (or phase modulator, PM), which prepares the PCB. The results obtained in our numerical experiments demonstrate that an appropriately chosen PM can significantly suppress the SI. In particular, one can try to suppress the effect of wandering by using a PM that deflects the laser beam as a whole in different directions, e_m , for fixed l . (This beam is a particular case of a PCB, as its phase changes in the plane $z = 0$. At the same time, its wave front remains flat for each realization of the PM in the plane transverse to the direction of propagation.) In this case, some of the directions, \vec{e}_m , will contribute to the compensation of wandering (see the insert in Fig. 2), and an average over the PM realizations signal at the detector,

$$\hat{I} = \frac{1}{M} \sum_{m=1}^M I_m^l,$$

can become rather stable-independent (or weakly dependent) of the atmospheric realization, l .

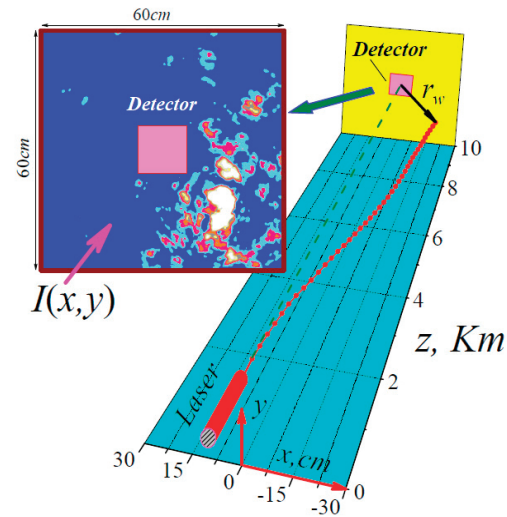


Fig. 1. Propagation of initially coherent laser beam through turbulent atmosphere ($\lambda = 1.55 \mu\text{m}$; radius of laser beam, $r_0 = 2\text{cm}$; $C_n^2 = 2.5 \times 10^{-14} \text{m}^{-2/3}$). An example of numerical simulations that demonstrates both fragmentation and wandering for a single atmospheric realization, l .

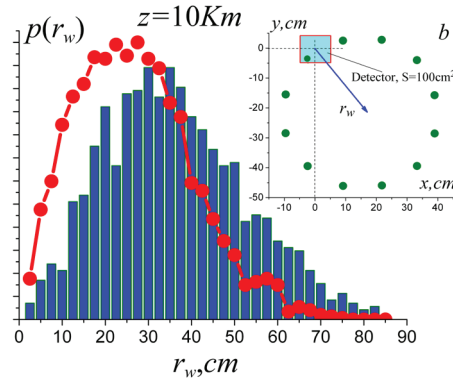


Fig. 2. The red curve is the distribution function, $p(r_w)$, for a coherent beam propagating initially along the z -axis, at $z = D = 10\text{ km}$. The insert shows the coordinates of the centers of PCB for random realization of atmosphere, $I_m^l(z) = \int \vec{r} I_m^l(x, y, z) dx dy / \int I_m^l(x, y, z) dx dy$, with angles, $\theta \approx \text{tg} \theta = r_{w, \max} / D$, to the axis z ($p(r_{w, \max} = 25\text{ cm}) = p_{\max}$). The blue color diagram, $p(r_c)$, corresponds to a distribution function for a PCB.

In [8] we present the results of the numerical experiment for a fixed set of vectors, $\vec{e}_m^{(i)}$, $m = 1, 2, \dots, 12$; $i = 1, 2, 3$. These vectors are deflected from the z -direction by the same angles, $\theta_i = \rho_i / D$ ($\rho_{1,2,3} = 15\text{ cm}, 25\text{ cm}, 35\text{ cm}$), and their projections on the x - y -plane are homogeneously distributed azimuthally (the angle, φ , between the neighboring projections was chosen to be 30°). The angles θ_i were chosen by the following protocol. In Fig. 2, the distribution function, $p(r_w)$, is presented in a red color in the x - y -plane and at $z = D = 10\text{ km}$ for the coherent beam oriented initially along the z -axis. As one can see, the most probable deviation (wandering) is, $15\text{ cm} \leq r_w \leq 35\text{ cm}$. Then, for a PCB the angles θ_i were chosen by taking into account this wandering parameter.

Figure 3 demonstrates the main result, which indicates an efficient suppression of the SI derived by averaging over 36 directions of the PCB. Note that in this protocol, the set of realizations for the PM can be chosen to be small, and it can be similar for all atmospheric realizations: A random choice of realizations from a continuous set, $\{\vec{e}\}$, is not required.

In conclusion, our numerical results demonstrate that the intensity scintillations can be significantly reduced by using a partially coherent beam with a particular realization of the phase modulator. This direction appears to be very promising for real implementations.

For more information contact Gennady Berman at gpb@lanl.gov.

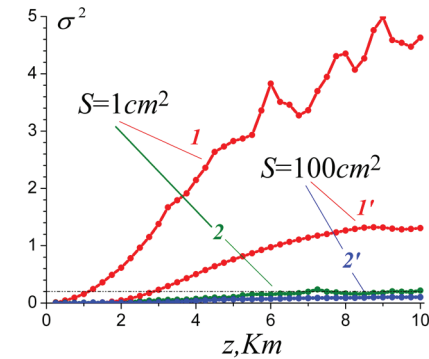


Fig. 3. The dependence of $\sigma^2(z)$ on distance (curves 1, 1' the result for a coherent beam); S is the area of detector. Curves 2 and 2' are calculated for average signal

$$\hat{I} = \frac{1}{36} \sum_{m,i} I_{m,i}^l.$$

- [1] L.C. Andrews, R.L. Phillips, C.Y. Hopen, *Laser Beam Scintillation with Applications*, SPIE Press (2001).
- [2] J. Pan et al., in *Wireless and Mobile Communications*, Proc. SPIE 4911, 58 (2002).
- [3] G.P. Berman, A.A. Chumak, A.A., V.N. Gorshkov, *Phys. Rev. E* **76**, 056606 (2007).
- [4] G.P. Berman et al., *Opt. Commun.* **280**, 264 (2006).
- [5] G.P. Berman, A.A. Chumak, *Phys. Rev. A* **74**, 013805 (2006).
- [6] G.P. Berman, A.R. Bishop, D.C. Nguyen, B.M. Chernobrod, V.N. Gorshkov, US Patent 7,603,038 (2009).
- [7] G.P. Berman, A.A. Chumak, *Phys. Rev. A* **79**, 063848 (2009).
- [8] G.P. Berman et al., SPIE Photonics West, (2010).

Funding Acknowledgments

Office of Naval Research (ONR)

Adaptive Finite Element Methods for Turbulent and Reactive Flow

David Carrington, David Torres, T-3

We have ongoing efforts to provide accurate solutions for time-dependent (unsteady) heat transfer [1], geophysical modeling (including atmospheric transport processes), indoor environmental transport [2], and, most recently, multiphase reactive flow using adaptive Finite Element Methods (FEM). For the last item, a method particularly well suited to combustion simulation is being developed in both 2D and 3D.

The fundamental mathematical model for a fluid momentum solution is based on predictor-corrector equation splitting. By using the pressure (density) solution, the predicted velocity or momentum is projected to the correct state. Solutions to energy and turbulent closure models are performed by scalar transport equations. The time-dependent solution is currently evolved by a backward Euler method, although second-order-in-time evolutions are possible. The resulting system of equations, particularly the Poisson equation for pressure or density, is solved with a preconditioned conjugate-gradient (PCG) solver package [3] where matrix preconditioning comes from in situ stationary methods [4].

Grids and Automatic Mesh Refinement. Simulation is accomplished by representing a domain in a discrete manner. FEM allows for the use of nonorthogonal or unstructured grids. A complex domain is represented with various discrete type elements, typically hexahedrals and trapezoids in

3D models. Grids for complex domains such as atmospheric boundary layers and internal combustion engines are shown in Fig. 1.

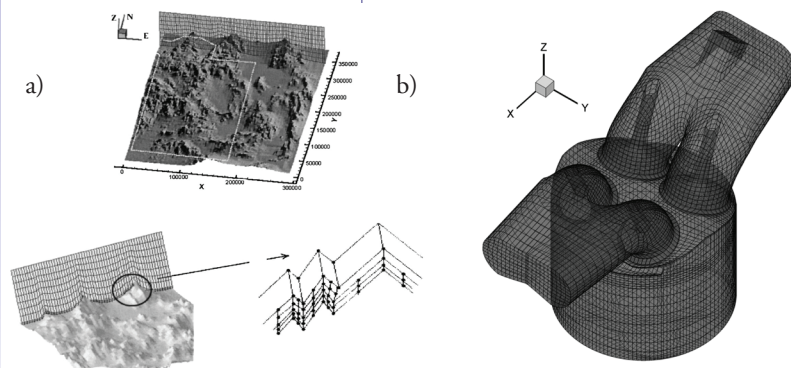
The use of h -adaptation, or mesh refinement, yields accurate solutions and rapid convergence rates. Exponential convergence occurs when higher order interpolation is employed, creating a hp-adaptive method [7]. The idea of using self-adaptive grids became apparent around 1970 – practically solving large-scale problems requires an optimal grid. The goal of adaptation is to achieve the proper mesh, one that significantly improves the computational speed and resolution, and reduces error. Embedding elements within the existing grid is a form of h -adaptation, a quickly processed procedure. There are many error estimators that can be used to control the adaptive process, e.g., the element-residual method, interpolation methods, and stress-error methods. The use of a simple residual-error estimator as described by Ainsworth and Oden [8] is used in the examples below.

For turbulent incompressible flows, an Euler-Lagrange variational formulation is developed resulting in a fractional-step method [9]. This method is applicable to the turbulent heat transfer of incompressible flows, atmospheric dispersion modeling, indoor air pollution modeling, solidification, and non-Newtonian incompressible flows. The model is also well suited for determining the wind's power over complex terrain – useful for precisely locating wind turbines.

For turbulent combustion modeling a characteristic-based split (CBS) is developed as described by Zienkiewicz and Codina [10]. The model spans all flow regimes and is applicable to the above-mentioned modeling areas in addition to compressible flow in the subsonic, transonic, and supersonic regimes. A second-gradient shock capturing scheme is employed along with h -adaptation to capture shocks – the specifics of the general method can be found in a recent report [11].

The backward-facing step is a benchmark problem indicative of how implemented algorithms/models will perform for many important aspects of flow. Study of flow and heat transfer over a backward-facing step with its richness of physics and availability of experimental data provides a rigorous test of the code. Flow phenomena

Fig. 1. a) Initial grids for atmospheric boundary layer [5], and b) unstructured grid for internal combustion engine for KIVA-4mpi simulation [6].



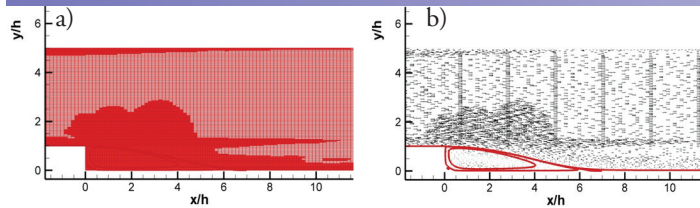


Fig. 2. a) Two levels of mesh enrichment, and b) vectors and streamlines on refined grid.

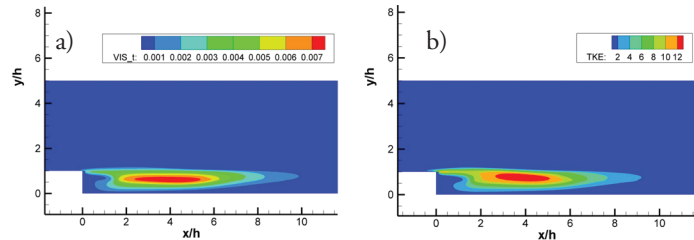


Fig. 3. a) Effective viscosity behind step, and b) turbulent kinetic energy behind step.

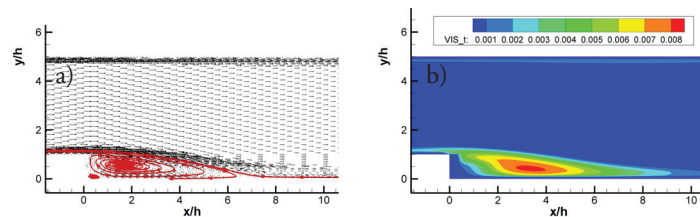


Fig. 4. a) Vectors and streamlines on refined grid, and b) effective viscosity behind step.

include unsteady behavior, separation, recirculation, and reattachment. Vogel and Eaton's experiments [12] provide experimental data for convective turbulent flow over a backward-facing step having an expansion ratio of 1.25.

Shown in Fig. 2 are results from a k - ω turbulence closure model using the projection method for the experimental setup of Vogel and Eaton with a Reynolds number (Re) = 28,000 [13]. Figure 2a shows the final grid, where the initial grid was enriched once nearly everywhere. More cell enrichment was needed as determined by a residual-error estimator in the boundary layer and near the expansion. Figure 2b shows the velocity vectors and streamlines behind

the expansion. The recirculation region size matches experimental results, at $6.4h$, where h is the step height.

Figure 3 depicts the effective viscosity and turbulent kinetic energy behind the step, respectively. The contours are in general agreement with other results. The results compare favorably with Vogel and Eaton's findings.

Turbulent Reactive Flow – Combustion Modeling. Resolution ushers in a need for greater computational power – a need for parallel processing. But solving the resolved problem with more computing power alone is not the only development required for achieving a desired accuracy. There are limits to what can be efficiently achieved by decomposing domains with moving parts and distributing the processing. It is generally better to have algorithms that are more accurate and also provide for higher resolution and accuracy only where and when it is required in the simulation process.

The KIVA team is facilitating combustion engine modeling by improving accuracy and robustness of the algorithms. To address the needs of the combustion engine community, including universities and industry, we recently began developing the CBS method. The method, combined with higher-order polynomial approximation (p -adaption) and h -adaption to form an hp -adaptive method, is being developed to provide for accurate and robust solutions in the next generation of KIVA software [11]. Shown in Fig. 4 is a solution using the CBS FEM method for the Vogel and

Eaton experiment previously described. At $Re=28,000$ the inflow is 17m/s (Mach number ~ 0.05), and the results compare favorably with other solutions and data. This is near the flow velocity at a cylinder's intake for a typical internal combustion engine running at 1000 rpm.

For more information contact
David Carrington at
dcarrington@lanl.gov.

- [1] D.B. Carrington, LA-UR-06-8432 (2007).
- [2] D.W. Pepper, D.B. Carrington, *Indoor Air Pollution Modeling*, Imperial College Press, **360** (2009).
- [3] W.D. Joubert, G.F. Cary, *Proc. CO Conf. Iterative Methods*, **1**, (1994).
- [4] D.B. Carrington, *Numerical Heat Transfer, Part B: Fundamentals* **53**, 1 (2008).
- [5] D.B. Carrington, D.W. Pepper, *Communications in Numerical Methods in Engineering* **18**, 195 (2002).
- [6] D.J. Torres, Y. Li, S.C. Kong, *Computers and Fluids*, **39.2**, 302 (2010).
- [7] X. Wang, D.B. Carrington, D.W. Pepper, *Computational Thermal Sciences*, **1**, 121 (2009).
- [8] M. Ainsworth, J.T. Oden, *Posteriori Error Estimation in Finite Element Analysis*, John Wiley and Sons, Inc., **20** (2000).
- [9] D.B. Carrington, D.W. Pepper, *Numerical Heat Transfer, Part A* **41**:6-7, 555 (2002).
- [10] O.C. Zienkiewicz, R. Codina, *Intl J. Numerical Methods Fluids*, **20**, 869 (1995).
- [11] D.B. Carrington, LA-UR-09-06527 (2009).
- [12] J.C. Vogel, J.K. Eaton, *Journal of Heat Transfer*, **107**, 922 (1985).
- [13] D.B. Carrington, X. Wang, D.W. Pepper LA-UR-09-7843 (2009), to be published in: *ICCES'10: International Conference on Computational & Experimental Engineering and Sciences* (2010).

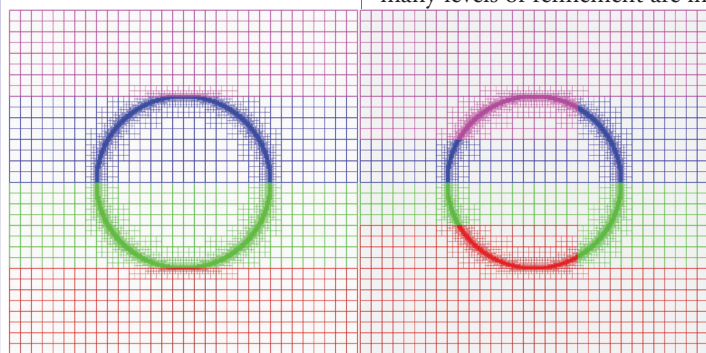
Funding Acknowledgements

DOE Energy Efficiency and Renewable Energy (EERE), Advanced Combustion Engine Research and Development

Development of a New Adaptive Mesh Refinement Framework for Multiphysics Simulations

William W. Dai, HPC-5

Fig. 1. The distribution of cells among four processors in a mesh with ten levels before (left image) and after (right image) the load balance. Each color represents the cells for which one PE is responsible. Before the balance, two middle processors each have about 0.6 million cells, whereas the other two processors each have about 6000 cells. After the balance, the middle two processors each have about 305,000 cells, and other two processors each have about 294,000 cells.



Adaptive mesh refinement (AMR) in analytical computation is the selective development of grids within grids where more detail is required than available using the parameters of the standard grid coordinates. There are three kinds of AMR. (1) Block-based AMR refines a predefined block when any cell within the block is marked to be refined. The advantages of the block-based AMR include the nature of structured meshes of each block and relatively simple data structures. But it sometimes over-refines meshes. (2) Cell-based AMR [1] refines only those cells that are supposed to be refined. But in cell-based AMR, cells, including the cells that are not refined, are typically treated cell by cell, even for structured meshes. Therefore cell-based AMR often loses the advantage of the nature of structured meshes. (3) Patch-based AMR dynamically groups to-be-refined cells to form rectangular patches and thus it combines the advantages of block- and cell-based AMR, i.e., the nature of structured meshes and the sharp regions of refinement. But patch-based AMR has its own difficulties. For example, existing patch-based AMR typically cannot preserve symmetries of physics problems. Since patches in patch-based AMR are dynamically generated and have different sizes, efficiently managing patches and communications in parallel environments presents a challenge when many levels of refinement are involved.

In this work, we have developed a patch-based AMR framework through which refined meshes are able to preserve symmetries of physics problems. Through the framework, the grid cells are refined near shock fronts, shear layers, and material interfaces. The development of the framework has been focused on its use

for multiphysics simulations and compressed material data structures. The approach for AMR consists of clustering, symmetry preserving, mesh continuity, flux correction, communications, management of patches, and dynamic load balance. The special features of this patch-based AMR include symmetry preservation, efficiency of refinement across shock fronts and material interfaces, special implementation of flux correction, and patch management in parallel computing environments.

For symmetry, if to-be-refined cells are symmetrical with respect to the x- or y- or z-axis, we want the refined cells determined through clustering algorithms to preserve the symmetry. We have developed additional steps, on top of a k-means algorithm developed by MacQueen [2], to preserve the symmetries. We have also implemented a procedure to ensure that the resolution in the resulting mesh through multiple levels of refinement smoothly changes from location to location. If a grid cell undergoes m levels of refinement, its immediate neighboring cells must undergo at least $(m-1)$ levels of refinement.

One of the advantages of patch-based AMR is that physics solvers are implemented on uniformly spaced structured meshes on each patch. To do that, the values of physics variables on the ghost cells of patches have to be obtained before physics solvers can be executed. We have implemented a procedure in which each computer processor communicates with each of its neighboring processors only once for

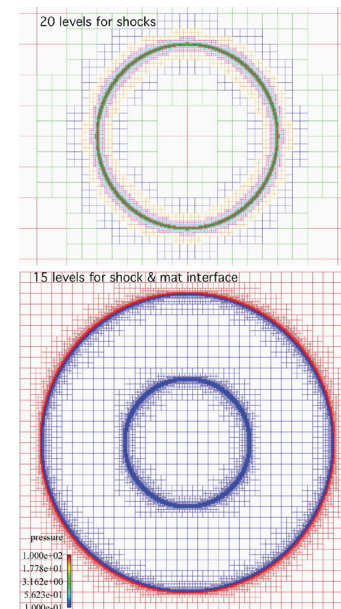


Fig. 2. The top mesh has 20 levels of cells near a circular shock front, and the bottom mesh has 16 levels of cell near a shock front (the outer circle) and a material interface (the inner circle). The cells of each level in the top image are shown through one color.

the values of ghost cells during one time step, no matter how many levels of refinement or which physics solvers are involved. Typically, the thickness of the layer of ghost cells depends on the physics solvers to be used, and the higher-order a solver is, the thicker the layer of ghost cells. In our framework, the thickness of the layer is a parameter, which can be set to the one appropriate for the physics solvers.

The number of patches in patch-based AMR may dramatically increase with the size of physics problems and the number of levels of refinement. The management of these patches may significantly influence the performance of AMR in parallel environments. We have developed an approach in which each processor manages only the minimum number of patches needed for updating physics variables on the processor. This approach yields far superior performance.

In a parallel computing environment, each computer processor is responsible for a certain amount of workload. Because of the dynamic nature of shock fronts and material interfaces, the workload of each processor also dynamically changes. We have implemented an approach for dynamic load balance. The approach is not for redistribution of workload but for the minimum data movement based on the existing workload and the connectivity of the subdomain each processor is responsible for. Figure 1 shows the distributions of grid cells in a mesh of 10 levels among four processors before and after the load balance. In the figure each color is the region for which one processor is responsible. Before the load balance, two middle processors each have about 0.6 million cells, whereas two other processors each have about 6000 cells. After the load balance, the middle two processors each have about 305,000 cells, and of the other two processors have about 294,000 cells.

Numerical Examples. Typically we refine shock fronts, shear layers, and material interfaces. For shocks and shear layers we use shock jump conditions and derivatives as the criterion for refinement. For material interfaces, we refine any cell with more than one material or mixed cell, and any pure cell whose immediately neighboring cells contain other materials. Figure 2 shows two meshes, the top one with 20 levels of cells near a circular shock front, and the bottom one with 16 levels of cells near a shock front (the outer circle) and

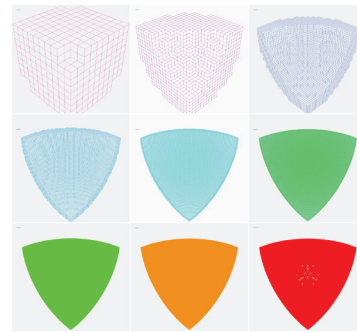


Figure 3. One eighth of a mesh with 10 levels of cells near a spherical shock front. The cells at each level are supposed to be buried in the cells of the previous level.

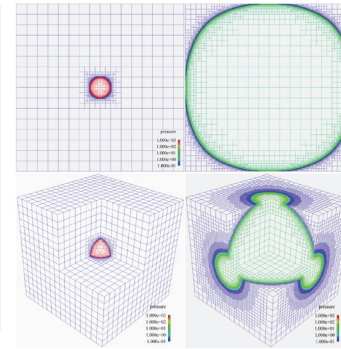


Figure 4. Distributions of pressure in 2D and 3D simulations at the initial time (the left images) and when the circular and spherical shocks reach the boundaries of the simulation domains. The refinement follows shocks propagating outward.

a material interface (the inner circle). Each color in the top image represents one level of cells. Only the first few levels of coarse cells are recognizable in the figure. The numbers of the 20 levels of cells in the top mesh, from the coarsest to finest, are 256, 304, 592, 1166, 2320, 4624, 9232, 18448, 36880, 73744, 147440, 294864, 589520, 1178128, 2352432, 4690128, 9321936, 18406656, 35890816, and 68047520, respectively. If one cell of a base patch were fully refined 19 times, the number of the finest cells in the cell would be 274,877,906,944.

The framework also works for the 3D case. Figure 3 shows one-eighth of refined cells near a spherical shock front with 10 levels of cells.

Cells at each level are actually buried in the cells of the previous level. The white dots in the last image in the figure indicate that the cells at the second-finest level do not undergo the refinement at these spots.

[1] M. Gittings et al., *Comput. Sci. Disc.* **1**, (2008); 015005, doi:10.1088/1749-4699/1/1/015005.

[2] J.B. MacQueen, "Some methods for classification and analysis of multivariate observations," *Proc. 5th Symp. Math. Stat. Probab.*, University of California Press (1967).

Funding Acknowledgments

DOE, NNSA Advanced Simulation and Computing Program Radiation Hydrodynamics on Roadrunner Project and Productivity Project

This AMR framework has been applied to some simple hydrodynamics calculations. Figure 4 shows the distributions of pressure in 2D and 3D simulations at the initial time (the left images), and when the circular and spherical shocks reach the boundaries of the simulation domain. The refinement follows the shock propagating outward.

We will also work on physics solvers involving multimaterials and compressed material data structures. When tens of materials are involved in large-scale simulations, compressed material data structures become very important.

For more information contact William W. Dai at dai@lanl.gov.

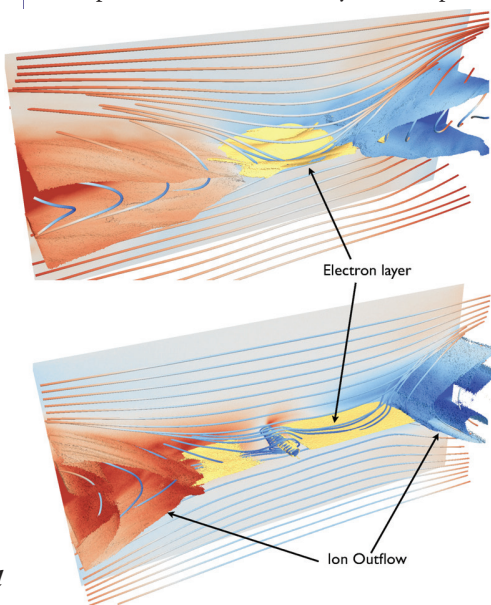
VPIC—Magnetic Reconnection

William S. Daughton, XCP-6, Vadim Roytershteyn, T-5; Lin Yin, Brian J. Albright, XCP-6; Ben Bergen, CCS-2; Kevin J. Bowers, XCP-6

Magnetic reconnection is a basic plasma process involving the rapid conversion of magnetic field energy into various forms of plasma kinetic energy, including high-speed flows, thermal heating, and highly energetic particles. These dynamical features are usually associated with changes in magnetic topology, which are conveniently viewed in terms of the breaking and reconnection of magnetic field lines. This process of *magnetic reconnection* is thought to play an important role in a diverse range of applications including solar flares, geomagnetic substorms, magnetic confinement fusion devices, and a wide variety of astrophysical problems.

In many applications, magnetic reconnection couples the large-scale plasma motion to fast dynamical processes at microscopic

Fig. 1. Open boundary simulations for neutral sheet geometry feature two types of secondary instabilities within the electron layer: an electromagnetic kink wave (top) and flux rope formation (bottom). The boundary conditions permit inflow of new plasma (the top and bottom) and outflow of reconnected plasma (left and right). The outflow jets are visualized with a particle density isosurface colored by ion outflow velocity, while the central electron current sheet (yellow) corresponds to an isosurface of electron current density. Sample magnetic field lines are colored by the magnitude of the magnetic field. Simulation was performed with mass ratio $m_i/m_e = 300$, using 4096 ranks, 245 million cells, and 147 billion particles.



scales. Many of the outstanding scientific challenges are related to this coupling between vastly different spatial and temporal scales.

High-temperature plasmas are very good electrical conductors, which implies that the magnetic field is constrained to move together with the plasma. For reconnection to proceed, it is necessary to break

this constraint within so-called *diffusion regions*, which are quite small in comparison to the macroscopic scale but play a critical role in the evolution. Researchers are working to understand the basic physics of these regions, as well as the coupling to the larger system. Due to the complexity of the relevant nonlinear processes, simulations have played an important role in this scientific progress. Most previous studies have focused on 2D models using a variety of fluid and kinetic descriptions. With increasing computer power, these simulations have progressed to larger systems and raised new questions. In particular, recent 2D kinetic simulations have demonstrated that diffusion regions often form elongated current sheets that are unstable to various plasma instabilities [1]. Understanding the 3D evolution of these layers is a formidable challenge that requires petascale computing.

To address these questions, LANL scientists are utilizing the 3D kinetic plasma simulation code VPIC [2], which provides a first-principles description of the relevant physics. The focus is to better understand the role of plasma instabilities on the 3D evolution of reconnection layers in both space and laboratory plasmas. The simulations performed on Roadrunner were of unprecedented scale and complexity, using upwards of 4096 ranks, ~200 billion particles, and requiring careful attention to boundary conditions, collisional physics, and new diagnostics. Despite these complications, the simulations achieved a factor of ~3x speedup using the IBM Cell Broadband Engine (Cell BE) processor. Here, we highlight a few science results emerging from these efforts.

In space and astrophysical applications, reconnection typically occurs in collisionless regimes and forms structures on both ion and electron kinetic scales. Since the macroscopic scales are vastly larger, it is not possible to study the global evolution while simultaneously resolving these small scales. To make progress, it is necessary to reduce the scale separation by employing artificial ion to electron mass ratios in the range $m_i/m_e \sim 100\text{--}300$, and to focus on thin current layers, which are the preferred sites for the onset of reconnection. Open boundary conditions [1] were employed to allow reconnection to develop over longer times, effectively mimicking the dynamics of a larger system using relatively small

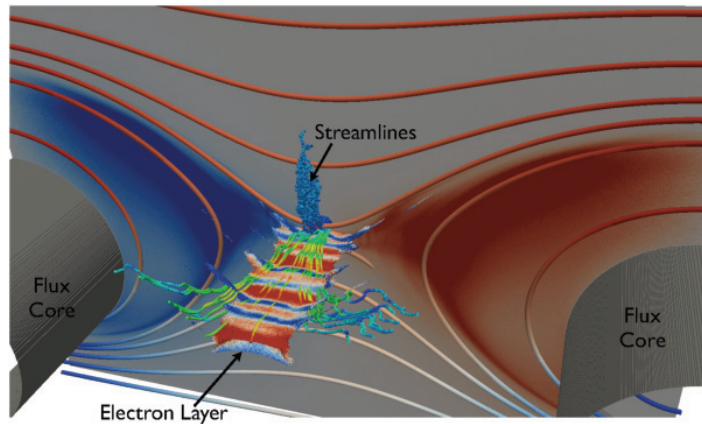


Fig. 2. Simulations were performed on Roadrunner with geometry and boundary conditions relevant to the MRX experiment [4,6] including the influence of weak Coulomb collisions [7]. The reconnection process is driven by reducing the currents inside the flux cores (grey cylinders), which pulls magnetic flux inward towards the cores. The resulting ion flow velocity is illustrated on the back cutting plane along with characteristic magnetic field lines. The central electron current sheet is illustrated with an isosurface of the current density, colored by a vertical component of the current density to show the plasma instability within the electron layer. Sample electron streamlines illustrate the electron flow from the upstream region through electron layer into the downstream region. This simulation was performed with mass ratio $m_i/m_e = 300$ using 2880 ranks, 720 million cells, and 144 billion particles.

simulation domains. Several of these simulations focused on neutral sheets, where the initial magnetic field reverses sign across the current layer and goes to zero in the center. After the onset and initial evolution, the diffusion region features highly elongated electron scale layers that are unstable to several plasma instabilities. As illustrated in Fig. 1, these modes include an electromagnetic wave that leads to *kinking* of the electron layer, and a secondary reconnection instability that gives rise to flux rope formation. The kink wave is similar to recent predictions from kinetic theory [3], while the flux ropes are the 3D analogue of secondary islands previously reported in 2D simulations [1]. Although these processes are qualitatively similar to recent 3D electron-positron ($m_i = m_e$) simulations [4], this is the first time they have been observed for the high mass ratio limit $m_i/m_e = 300$ that is now feasible with

Roadrunner. These new results are of great interest from several perspectives. First, the kink wave can potentially give rise to a wave-induced dissipation, which may dramatically alter the reconnection process. In addition, the formation of flux ropes introduces a strong perturbation of the diffusion region topology, modifying both the size and internal structure. These modifications can influence the macroscopically important parameters, such as the rate of plasma flow into the layer.

Magnetic reconnection is also of great interest in laboratory experiments, which offer the capability to study the structure and dynamics in a controlled setting. Three-dimensional kinetic simulations may serve as a bridge to help extrapolate ideas that have been validated in the laboratory to regimes of direct relevance to space and astrophysical plasmas. The Magnetic Reconnection eXperiment (MRX) at the Princeton Plasma Physics Laboratory (PPPL) is one leading experiment that has reported detailed measurements of the diffusion region, including the kinetic structure of the electron layer [5]. Weak binary Coulomb collisions are necessary to properly describe these plasmas, but the applicability of fluid models is questionable. The kinetic approach within VPIC includes experimental boundary conditions appropriate to MRX [6] and a Monte Carlo treatment of the collisions [7]. On Roadrunner, a series of simulations were performed to examine the influence of collisionality and plasma instabilities on the structure of the diffusion region. An example 3D simulation in Fig. 2 illustrates an electromagnetic instability within the electron current layer. This instability has certain similarities

with the kink instability discussed in Fig. 1. Researchers at LANL are working with the scientists at PPPL to perform comparisons with these new simulation results. These validation efforts will focus on the geometry of the electron layer and the observed electromagnetic wave spectra.

Together with ongoing theoretical work, these efforts are expected to shed new light on the influence of plasma instabilities on magnetic reconnection.

For more information contact William Daughton at daughton@lanl.gov.

- [1] W. Daughton et al., *Phys. Plasmas* **13**, 072101 (2006).
- [2] K. Bowers et al., *Phys. Plasmas* **15**, 055703 (2008)
- [3] W. Daughton, *Phys. Plasmas* **10**, 3103 (2003).
- [4] L. Yin et al., *Phys. Rev. Lett.* **101**, 125001 (2008).
- [5] Y. Ren et al., *Phys. Plasmas* **15**, 082113 (2008).
- [6] S. Dorfman et al., *Phys. Plasmas* **15**, 102107, 2008
- [7] W. Daughton et al., *Phys. Plasmas* **16**, 072117 (2009).

Funding Acknowledgments

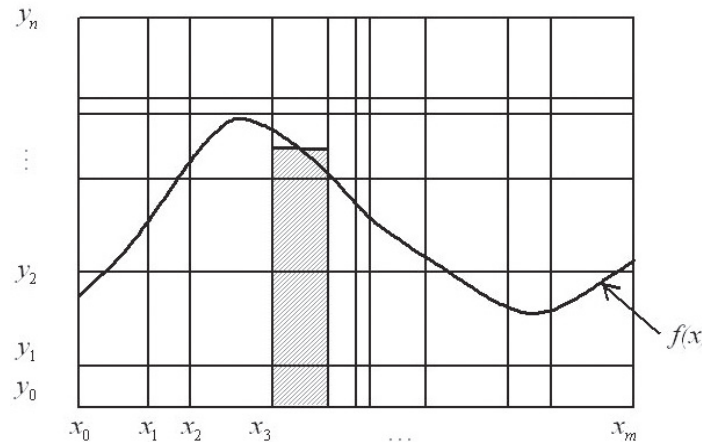
- LANL Research and Development Program
- DOE, NNSA Advanced Simulation and Computing (ASC) Program

Accurate Estimation of Geometrical Properties of a Volume-of-Fluid Interface on Nonuniform Rectangular Grids

Marianne M. Francois, CCS-2; Blair K. Swartz, T-5

Estimating the local curvature of an interface involves the local determination of normals (lines perpendicular) to the interface, and the rates that they turn as one moves along the interface. This is challenging for volume-of-fluid type methods of interface approximation because the interface between materials is then specified by the relative amount it cuts off from the computational cells that it crosses (also referred to as volume fraction data), rather than by a discrete set of points lying on the interface itself. In our work [1], we have generalized a volume-of-fluid method, called the height function method, to nonuniform grids of rectangular cells (see Fig. 1). First, for each vertical column of mesh cells, the volume fractions determine an (integral) mean value (or column height) for the portion of the interfacial curve that crosses that column. We then demonstrate analytically and numerically that three successive

Fig. 1. Illustration of an irregular rectangular grid and a known (integral) mean value of the unknown interfacial curve between the sides of one of the grid's columns.



(adjacent) mean values suffice to estimate interfacial curvature to second-order accuracy, the interfacial slope to third-order accuracy, and the interfacial curve's location to fourth-order accuracy—each at its own special points as one crosses the three columns (see Fig. 2). Second-order accurate curvature can be estimated anywhere by linearly interpolating the second-order accurate curvatures between two successive special points. And finally, we have determined the special locations where the curvature can be estimated to fourth-order accuracy when using five successive mean values instead.

Underlying all this is a result about the accuracy of the j th-derivative of the k th-degree polynomial that interpolates (matches) a function F at $k + 1$ "stencil points" placed irregularly in an interval of width h . Namely, for all smooth-enough functions F and for k fixed and h getting small, there are $k + 1 - j$ special points in the interval at which the error in the polynomial's j th derivative is of order $O(h^{k+2-j})$. This is one order higher than the usual $O(h^{k+1-j})$ error holding over the whole interval for that derivative. Increasing a method's order of accuracy has consequences for computational efficiency. For example, the error in a third-order $O(h^3)$ accurate method goes to zero at least as fast as the third power h^3 of the interval width h . Then, cutting the width by a factor of 2 would cut the error by a factor of $2^3 = 8$.

These special points of accuracy in the interval are the $k + 1 - j$ (real) zeroes of certain F -independent polynomials, of degree $k + 1 - j$, with coefficients depending on the interval's stencil points. For example, take $j = k$, the polynomial's degree. It is classic that the interpolating polynomial's k th derivative is a first-order $O(h)$ accurate estimate of the k th derivative of F anywhere in the h -sized interval containing the stencil points. But it is second-order $O(h^2)$ accurate only at one special point; namely, at the average location of the $k + 1$ stencil points.

In the context of the first paragraph, suppose the (unknown) interfacial curve $f(x)$ crosses k adjacent columns of cells. The edges of the columns determine $k + 1$ successive points x_1, x_2, \dots, x_{k+1} . Let $F(x)$

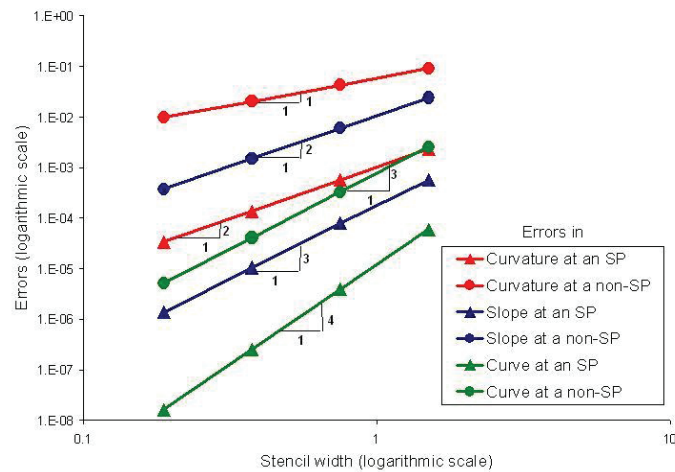


Figure 2. Log-log plots of the errors in interfacial curvature, interfacial slope or the interfacial curve's location—at an appropriate special point (SP) of accuracy or at a point that is not a special point of accuracy—using a cosine test function and a shrinking stencil of four nonuniformly distributed mesh points. The indicated slope of a plot gives the error's order of accuracy.

be an indefinite integral of the function $f(x)$. Then $F(x_i)$ at the $k + 1$ successive locations x_i is calculated using cumulative sums of the k known successive integral mean values of f , weighted by the successive distances between the x_i . The $k + 1$ points $(x_i, F(x_i))$ are now interpolated by a k th degree polynomial $(PF)(x)$. Its first derivative $(PF)^{(1)}$ approximates the unknown curve $f = (F^{(1)})$, while $(PF)^{(2)}$ and $(PF)^{(3)}$, respectively, approximate $f^{(1)}$, and $f^{(2)}$. Then the results in Fig. 2, for the curvature $f^{(2)}/(1+(f^{(1)})^2)^{3/2}$, the slope $f^{(1)}$, and the curve's location f , are obtained using $k = 3$ successive mean values and the derivatives $j = 3, 2, 1$. The $O(h^4)$ accurate curvature using five successive mean values involves $k = 5$ and $j = 3, 2$.

For more information contact Marianne Francois at
mmfran@lanl.gov.

[1] M.M. Francois, B.K. Swartz, *J. Comp. Phys.* **229**, 527 (2010).

Funding Acknowledgments

LANL Directed Research and Development Program—Exploratory Research (ER)

Conformal Refinement and Coarsening of Unstructured Quadrilateral Meshes

Rao V. Garimella, T-5

We have developed a novel method of adapting unstructured quadrilateral (quad) meshes for accurately capturing special features of unsteady solutions of partial differential equations, or PDEs. Adaptive refinement enables the capturing of complex solution features by focusing refinement in critical areas without having to refine the mesh everywhere. In our method, we can refine or coarsen unstructured quadrilateral meshes as driven by an error indicator or estimator while maintaining a conformal mesh (meaning no node of a refined quadrilateral lies inside an edge of a coarser quadrilateral). This implies that a numerical method that works on a fixed conformal mesh can be applied as-is to dynamically adapted meshes.

The adaptive mesh modification algorithm starts with tagging elements that must be refined because they do not adequately represent some geometric feature or because the solution error in these elements is too high. These elements and their edges are tagged for refinement (or coarsening), if necessary, to multiple levels below (or above) their current level of refinement. Once the appropriate elements have been tagged by the application, the mesh is coarsened

Fig. 1. Conformal refinement of triangles (left) and non-conformal refinement of quads (right).

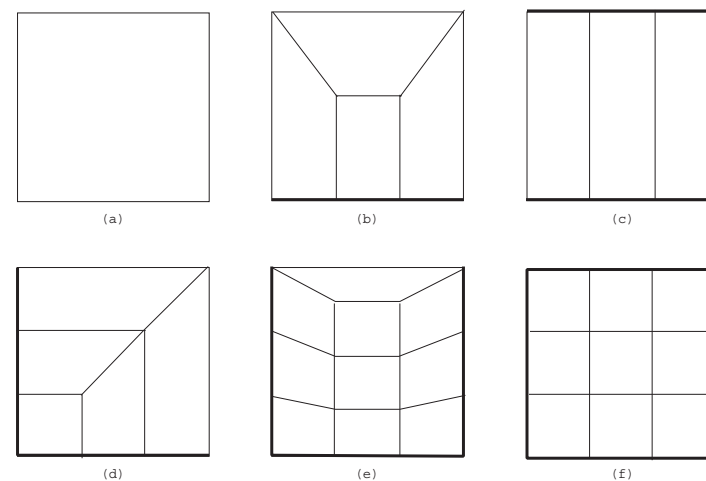
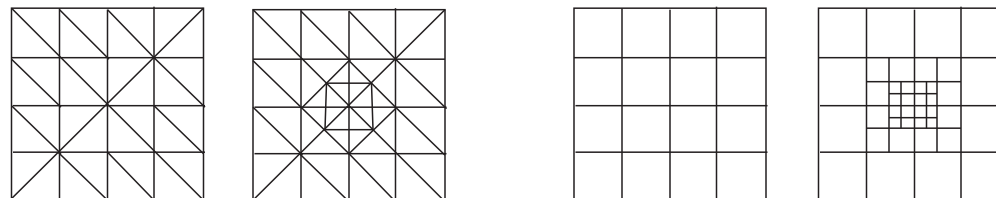


Fig. 2. Subdivision templates (thick edges are refined) (a) unrefined quad (b) one edge refined (c) two opposite edges refined (d) two adjacent edges refined (e) three edges refined (f) all edges refined (uniform refinement).

wherever the elements are smaller than they need to be. After coarsening, the mesh is refined wherever the elements are larger than they need to be. During both coarsening and refinement, the target refinement levels of elements are adjusted so that they are consistent with their siblings (children of their parents) and such that the target refinement levels of two adjacent elements do not differ by more than one. The one-level difference rule ensures that the number of templates required to make the mesh conforming is limited to a manageable number and that the mesh is smoothly graded.

Whenever a quad is uniformly refined, edges of adjacent quads are also refined. To make the mesh strictly conforming, these adjacent elements must also be subdivided into quadrilaterals such that there are no nonconforming nodes (Fig. 1). The templates used for subdividing quadrilaterals with different edges refined are shown in Fig. 2. Some of these templates have been described in previous works [1] and some are new.

The quads that result from uniform refinement of a parent quad are called regular elements, while quads resulting from refinement of one, two, or three edges of the parent quad are called irregular quads. An important aspect of this method is that irregular quads are never refined since their subdivision can result in substantial degradation of quality. Rather, if an irregular quad is considered to be too large, it is deleted along with its siblings, and the parent is refined to the required size.

First we show examples of a structured mesh and an unstructured mesh adapted to a superimposed line in the mesh (Fig 3). Any element that is intersected by the line is refined up to level 3 (level 0 is the original mesh). The superimposed line goes from $(-0.308207, 1.106007)$ to $(1.106007 - 0.308207)$ (Fig.4).

Next, two snapshots (Fig. 5) of a dynamic adaptation procedure are shown in which elements intersecting one of two expanding circles are refined to level 3 and elements intersecting both circles are refined to level 4. One circle is centered at $(0.0, 0.0)$, and the other circle is centered at $(1.0, 0.25)$. Both circles start with a radius of 0.11 with their radii increasing in increments of 0.05. As the circles grow, they intersect each other and eventually grow out of the domain.

For more information contact Rao V. Garimella at rao@lanl.gov.

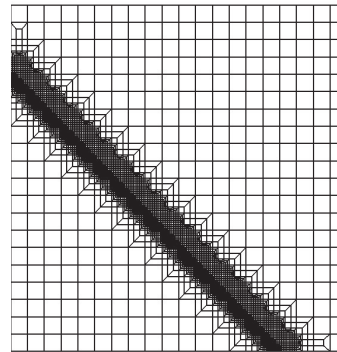


Fig. 3. A 20x20 structured mesh refined along a line.

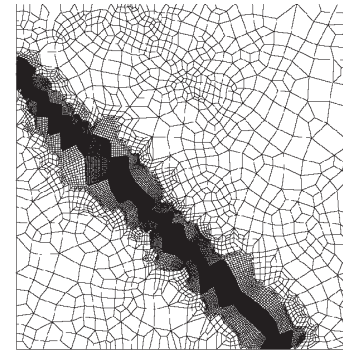


Fig. 4. An unstructured mesh refined along the same line.

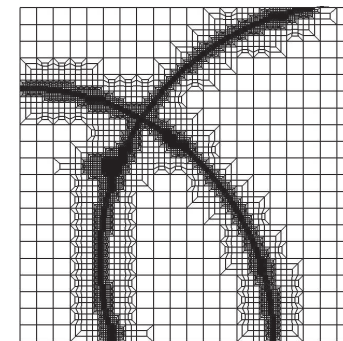
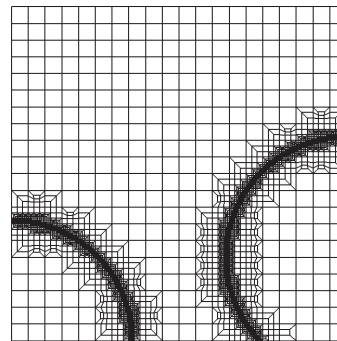


Fig. 5. Snapshots of dynamic mesh adaptation of a 20x20 structured mesh with respect to two expanding circles

[1] J.S. Sandhu et al., *Eng. Fract. Mech.* **50**, 727 (1995).

Funding Acknowledgments

LANL Directed Research and Development Program

Effective Shear Viscosity of Active Suspensions at Moderate Concentrations

Vitaliy Gyrya, Pennsylvania State University; Konstantin Lipnikov, T-5; Igor Aronson, Argonne National Laboratory; Leonid Berlyand, Pennsylvania State University

A number of recent experimental studies [1] have demonstrated that a suspension of self-propelled bacteria (*Bacillus subtilis*) at moderate concentrations may have an effective viscosity that is significantly (up to five to seven times) smaller than the viscosity of the ambient fluid. This is in sharp contrast to suspensions of hard passive inclusions, whose presence always increases viscosity. Detailed understanding of this viscosity phenomenon will help in the development of new materials and engineering solutions (e.g., improved micro-mixers). The 2D model developed in [2] captured an experimentally observed decrease of effective viscosity and provided an explanation for the underlying mechanisms.

The analysis was performed for moderate concentrations (about 9% by volume) of micro-swimmers that closely resemble the experimental settings, where a computer simulation is the only available tool. The analysis showed that the decrease in the effective viscosity observed in the physical experiments can be explained entirely from the point of view of hydrodynamic interactions. This is an important observation, since suspensions of bacteria represent a very complex system with a variety of phenomena occurring simultaneously, such as chemotaxis and secretion of proteins by living bacteria. These phenomena may influence the effective viscosity and are hard to isolate in physical experiments.

The key features leading to the decrease of viscosity are: 1) self-propulsion, 2) the elongated shape of the swimmers, and 3) the swimmer-swimmer interactions.

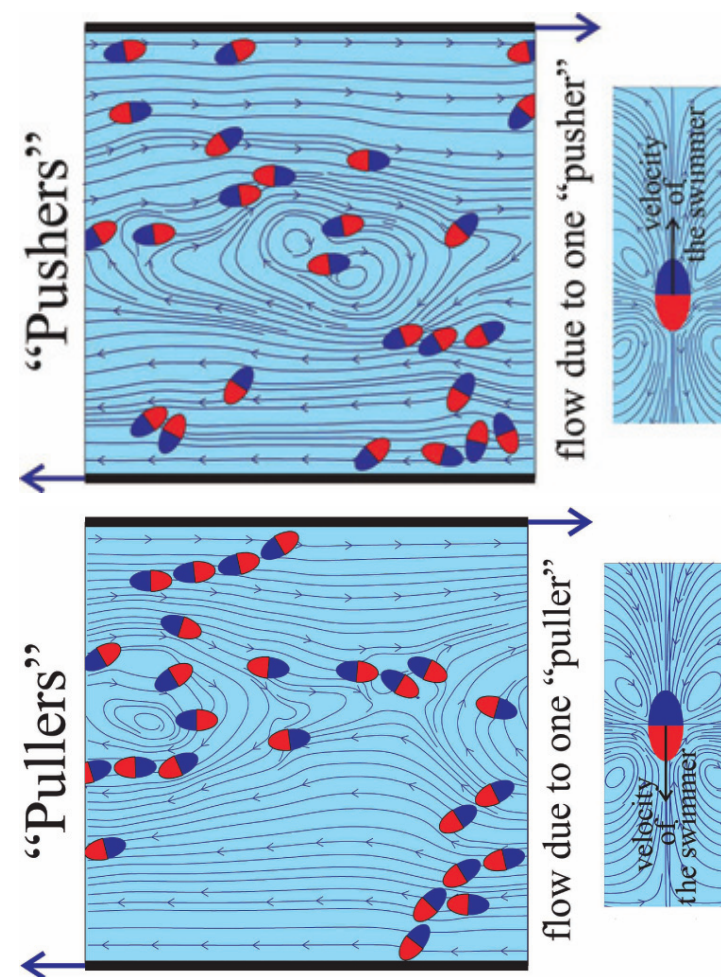


Fig. 1. Snapshots of micro-swimmers in the shear background flow. The blue half of the ellipse represents a solid surface, while the red half represents the surface covered by flagella. The top figure shows pushers ($f_p > 0$), which tend to swim side-by-side. The bottom figure shows pullers ($f_p < 0$), which tend to follow each other forming head-to-tail train-like structures.

The elongated body of the swimmer is modeled by an ellipse. The self-propulsion is modeled by distributing a force (modeling the action of the flagella) of magnitude f_p over the back half of the

ellipse. This force pushes fluid backward (away from the center of the swimmer), resulting in the forward propulsion of the swimmer. The other (front) half of the ellipse represents a solid surface (head of the bacterium). The numerical analysis and experimental observations show that it is important to distinguish between two types of swimmers: pushers ($f_p > 0$) – swimmers propelling themselves forward (e.g., *B. subtilis*), and pullers ($f_p < 0$) – swimmers propelling themselves backwards (e.g., some algae). Pushers tend to swim side-by-side, while pullers tend to follow each other, forming train-like structures (Fig.1).

The key step in the numerical analysis of the problem was the solution of the incompressible Stokes problem in the fluid domain with complex geometry and mixed boundary conditions. The solution was achieved using the mimetic finite difference (MFD) method [3]. The MFD method combines the mesh flexibility of the finite volume methods with the analytical power of finite element (FE) methods. The MFD method can be viewed as an extension of FE methods to unstructured polygonal meshes. The mesh flexibility simplifies mesh generation around swimmers that may have complicated shapes, e.g., the MFD method allows the mesh elements to have curved edges. Compared with FE methods, the MFD method minimizes the number of discrete unknowns (without loss of accuracy) by: 1) partitioning of the computation domain into a smaller number of elements that are polygons, and 2) using a smaller number of velocity and pressure unknowns only where they are needed for accuracy and stability of the discretization. The MFD method is second-order accurate (with respect to the local mesh size) for the velocity and first-order accurate for the pressure.

Numerically, the effective shear viscosity was measured by placing a suspension of micro-swimmers between two horizontal plates and forcing these plates to move in opposite directions with a constant velocity. The effective viscosity was defined as a coefficient of proportionality between the velocity of the plates and the force required to keep the plates moving.

The numerical simulations showed that the effective shear (Fig.2.) viscosity decays linearly as a function of the propulsion strength f_p , i.e., pushers decrease the effective viscosity. The linear trend continues for negative values of f_p , i.e., pullers increase the effective viscosity. For large values of the propulsion strength, the linear trend seems to change. We explain this by the finite size of the fluid domain in our simulations and the dynamics of a single swimmer in a shear flow.

For more information contact
Konstantin Lipnikov at
lipnikov@lanl.gov.

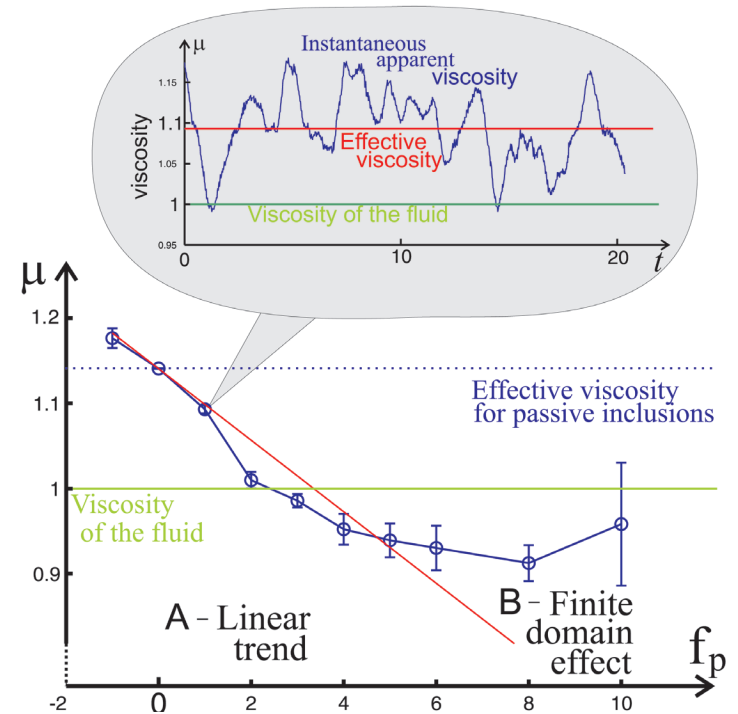


Fig. 2. Suspension of 25 swimmers at moderate concentration (9% by volume) in a unit square with shear velocity boundary conditions on top and bottom sides, and periodic boundary conditions on vertical sides. The green horizontal line indicates the viscosity of the ambient fluid. The blue oscillating line in the top picture shows instantaneous measurements of the shear viscosity. The red horizontal line indicates the time average of the instantaneous measurements. Linear reduction of the effective shear viscosity is observed when propulsion force f_p is less than 5. Wall effects become important for larger values of f_p .

- [1] A. Sokolov, I.S. Aronson, *Phys. Rev. Lett.* **103**, 148101 (2009).
- [2] V. Gyrya et al., LAUR 09-06018; (2009); submitted to *J. Math. Bio.*
- [3] L. Beirão da Veiga et al., *J. Comp. Phys.* **228**, 7215 (2009).

Funding Acknowledgments

DOE Office of Science Advanced Scientific Computing Research (ASCR) program in Applied Mathematical Sciences

High Atwood Number Rayleigh-Taylor Instability

Daniel Livescu, Mark R. Petersen, CCS-2; Steven L. Martin, Patrick S. McCormick, CCS-1

Molecular mixing as a consequence of stirring by turbulence is an important process in many practical applications. If the microscopic densities of the fluids participating in the mixing are very different, we refer to such flows as variable density (VD) flows in contrast to the Boussinesq approximation in which the densities are close. In VD flows, the velocity field is no longer solenoidal, even if the fluids participating in the flow are incompressible. VD mixing is encountered in atmospheric and ocean flows, astrophysical flows, combustion and many flows of chemical engineering interest. Many of these flows are driven by acceleration (e.g. gravity in geophysical and astrophysical flows) which, because the density is not uniform, leads to large differential fluid accelerations. If the acceleration is constant and the fluid configuration is unstable (i.e. density gradient points opposite to the body force), a fluid instability is generated in which small perturbations of the initial interface between the two fluids grow, interact nonlinearly, and lead to turbulence. This instability is known as the Rayleigh-Taylor (RT) instability and is of fundamental importance in a multitude of applications, from fluidized beds, oceans and atmosphere, to ICF and supernovae explosions. Due to RT instability, any perturbation with a wavelength larger than the cutoff due to surface tension (for the immiscible case) or mass diffusion (for the miscible case) will grow. As the perturbation grows, smaller and larger eddies are generated by nonlinear interactions and eventually the flow becomes turbulent.

In many cases, the density ratio between the two fluids is large (e.g. air interpenetrating helium has a density ratio of 7), yet most studies to date address the low density ratio case and no Direct Numerical Simulations have been performed for Atwood number, $A > 0.5$ (corresponding to a density ratio of 3), where $0 < A < 1$.

Previous results at Atwood numbers up to 0.5 [1,2,3] indicate that mixing becomes qualitatively different at high density ratios (the variable density case) compared to the case when the densities of the two fluids are commensurate (the Boussinesq approximation). Thus, in an idealized triply periodic configuration, [1] pure heavy fluid mixes more slowly than pure light fluid: an initially symmetric double delta density PDF is rapidly skewed. The density PDF skewness generation mechanism, $\langle \rho \nabla \rho \cdot \nabla \rho \rangle$, can be shown to be determined, [3] through changes in the magnitude of the density gradient, by the eigenvalues of the strain rate tensor and the relative alignment between the density gradient and the eigenvectors of the strain rate tensor. Both the eigenvalues and relative alignment are different in the pure heavy and light fluid regions. In the heavy fluid regions, the eigenvalues of the strain rate tensor have lower values, as the larger inertia of the heavy fluid leads to lower deformation rates of the fluid blobs. Larger inertia of the heavy fluid also leads to decreased alignment with the local flow. Thus, the local structure of the flow changes in response to the inertia of the fluid particles. Consequently, the inertia of the heavy fluid reduces the rate at which it is broken up by stirring, decreasing the local surface area of the pure heavy fluid blobs, which is related to the magnitude of the density gradient. The net result is that the magnitude of the density gradient is lower in the pure heavy fluid regions, along with the rate of molecular mixing, and the density PDF becomes skewed.

For the RT configuration at large density differences, this suggests that molecular mixing proceeds differently on the two sides of the RT layer. Experiments to date have not investigated this possibility. One consequence for the Rayleigh-Taylor case (as found from a 3072^3 simulation at $A=0.5$) [4] is that the penetration distance of the pure heavy fluid is larger than that of the pure light fluid. [2,3] The mixing asymmetry is likely also the cause of the bubble-spike anomaly (higher growth rate on the spike side compared to the bubble side), which was observed experimentally, but never explained.

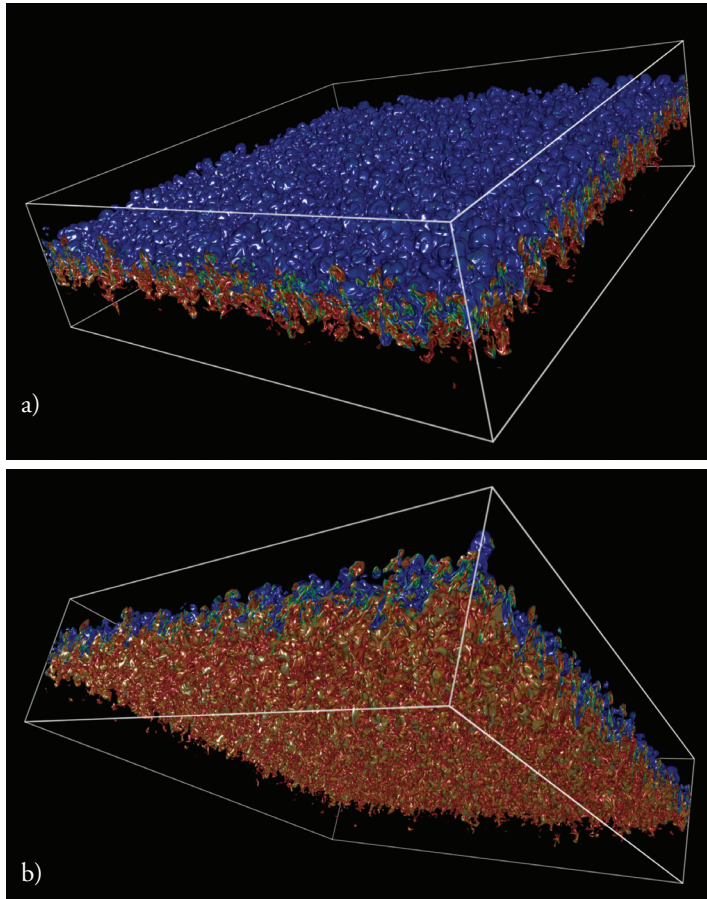


Fig. 1. (a) Three-dimensional visualization of the density field showing the asymmetry of the Rayleigh-Taylor mixing layer, with the development of "bubble" on the heavy fluid side (a) and "spikes" on the light fluid side (b).

As the Atwood number is increased above 0.5, the consequence of the mixing asymmetry becomes more obvious: the two edges of the mixing layer develop into "spikes" on the light fluid side and "bubbles" on the heavy fluid side. This can be clearly seen in Fig. 1, which shows the density field from a 2304×4096^2 Rayleigh-Taylor simulation at $A=0.75$, the largest fully resolved instability simulation to date. The simulation was performed on the ASC Dawn supercomputer, at Lawrence Livermore National Laboratory, during the open science runs. Most of the simulation was run on the full Dawn machine (147,456 processors). Currently, the simulation is being extended to 4608×4096^2 .

Three-dimensional visualization of such large data sets also imposes enormous challenges and novel techniques were developed for this. Thus, Fig. 1 was generated using a parallel, hardware-accelerated analytical ray casting algorithm.

-
- [1] D. Livescu, J.R. Ristorcelli, *J. Fluid Mech.* **605**, 145 (2008).
 - [2] D. Livescu et al., *J. Turb.* **10**, N13 (2009).
 - [3] D. Livescu, J.R. Ristorcelli, "Mixing Asymmetry in Variable Density Turbulence," in *Advances in Turbulence XII*, Springer (2009).
 - [4] W. H. Cabot and, A.W. Cook, *Nat. Phys.* **2**, 562 (2006).

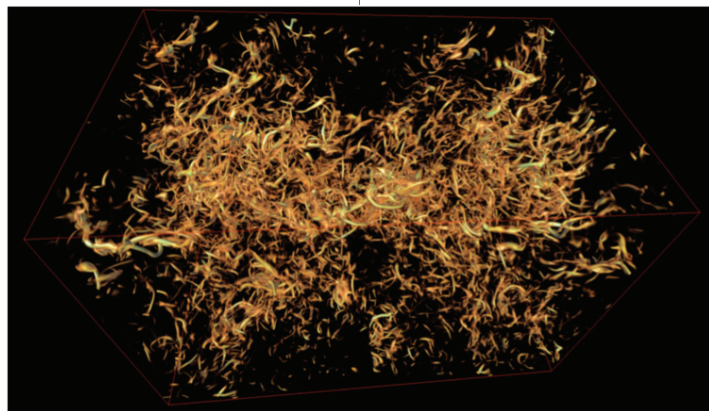
Funding Acknowledgments

LANL Directed Research and Development Program

Direct Numerical Simulations of Reacting Compressible Turbulence

Daniel Livescu, Jamaludin Mohd-Yusof, Timothy M. Kelley, CCS-2

Fig. 1. Enstrophy field in reacting compressible turbulence shows the rich phenomenology of the flame turbulence interaction. The very high resolution simulations highlight the presence of vortex tubes. Across the flame, fire polishing damps the smallest turbulence scales; however some of the vortex tubes are enhanced through baroclinic vorticity generation.



Although important progress has been made in recent years in our understanding of turbulence, complete quantification, description, prediction, simulation, and control still elude us. The problem is due in part to the very large range of spatio-temporal, dynamically relevant scales, but also to the multitude of problems that can be encompassed by the generic term turbulence. If ideal turbulence is in a homogeneous, isotropic Kolmogorov steady state, then nonideal turbulence can occur due to most practically relevant effects: time-dependence, anisotropy, inhomogeneity, coupling with active scalars, shock waves, exothermic reactions, etc. Unlike kinetic theory, where significant departure from a weakly perturbed local Maxwellian is exceptional, the analogous state of nonideal turbulence is what is typical, yet the only successful turbulence theory so far is Kolmogorov's 1941 theory. It is then not surprising that the turbulence theory is still centered on the ideas of Kolmogorov related to the existence of inertial range dynamics (e.g., a $-5/3$ range in the isotropic turbulence energy spectrum) and small scale universality. However, more and more evidence points to departures from universal laws in the energy spectrum due to intermittency and a

direct connection between the small and the (nonuniversal) large scales, especially in the presence of strong gradients [1] or buoyancy [2-4].

Most of the turbulence research to date has been concentrated on several canonical flows with periodic boundaries or simple jets, wakes, or boundary layers. Numerous modeling strategies have been proposed, and, while

there is no best strategy, each of the approaches has its own advantages and domains of applicability. Yet most practical flows are not canonical. In many situations, they are driven by acceleration as in ICF or cosmic explosions and may undergo exothermic reactions. In addition, radiation could have a significant effect, e.g., through heat gain or loss. For these complex flows, the limitations of the current modeling strategies, as well as the turbulence and mixing properties are unknown.

Turbulence theory and the subsequent model development rely on experimental or high resolution direct numerical simulations (DNS) data for development and verification and validation. This technique has emerged as a powerful research tool to study the physics of turbulence, verify and improve models, complement and even guide and help design better experiments [5]. The DNS technique seeks exact solutions of the governing equations, so that all relevant scales are accurately solved, using high resolution numerical simulations based on high-order accurate discretization algorithms. DNS rely on nondissipative high accuracy schemes and are conducted without resort to subgrid modeling or the introduction of artificial numerical dissipation or other algorithm stabilizing schemes. Such computations allow a degree of control in isolating specific physical phenomena that is typically inaccessible in experiments. With the recent advances in supercomputing technology and algorithms, it is now possible to perform simulations of simple flows at ranges of scales comparable to or even larger than in typical laboratory experiments. Petascale computing is expected to further increase the range of scales of the simulations and allow accurate calculations of more and more complex flows.

This study represents the first successful implementation of a large structured fluid dynamics code (CFDNS) on the IBM Cell Broadband Engine (Cell BE) architecture. The CFDNS code solves the compressible and incompressible Navier-Stokes equations in 3D using high-order compact finite differences or Fourier transforms in the periodic directions. Multiple species are allowed, each with realistic material properties and equations of state (EOS), as well as Cartesian, cylindrical, and spherical grid geometries. The serial speedup of the Cell BE version of the code is approximately 30x, which is

reasonable when considering the clock speed, parallelism and vectorization afforded by the Cell BE. Notably, the excellent performance of the individual memory controllers is responsible for this, since the low arithmetic intensity of the algorithm does not allow the actual compute power of the Cell's synergistic processing elements (SPEs) to be utilized. This serial speedup prompted us to perform significant modifications to the parallel code design, as indicated above, which lead to overall speedup in the range of 20x compared with the AMD Opteron-only version [6,7].

The model problem addressed by the study is the flame-turbulence interaction under the complex conditions characterizing the early stages of a type Ia supernova. These conditions are novel and have no direct analogue on earth. This makes them interesting for testing new physics, but it also means that our terrestrial intuition regarding flames can be misleading. For example, it is thought that laboratory flames in the $Ka \gg 1$ regime simply go out because they are unable to maintain their heat in the presence of so much turbulence. But the flame in a supernova can never go out until the star comes apart and, in terms of local flame variables, that takes a very long time. Although several mechanisms for detonation have been proposed, the debate around deflagration/detonation models is still not settled. Moreover, most turbulent flame simulations so far, under these conditions, are in the low Mach number approximation and no DNS have been performed for such flows.

Using the Roadrunner supercomputer, researchers have performed the largest reacting compressible turbulence simulations to date. The flow conditions considered for the reacting simulations correspond to well-stirred single component burning, $^{12}\text{C}(^{12}\text{C},\text{g})^{24}\text{Mg}$, relevant to Type Ia supernovae. The flame advances into the cold fuel (^{12}C) in a C-O mixture and leaves behind hot product (^{24}Mg). Inflow/outflow boundary conditions are imposed in the flame propagation direction. The physical transport properties are appropriate for the astrophysical situation investigated and are calculated within new modules added to the code. Thus, the thermal transport includes both radiative and electron transport (accounting for degenerate regimes) contributions. The equation of state considers radiative, ion, and electron contributions. To reduce the computational effort, precomputed tables for the

transport properties, EOS, and nuclear energy rates are used.

The simulations were performed in stages: a) generate 1D reacting flow profiles as initial conditions, b) generate inflow turbulence by performing triply periodic simulations with a background velocity matching the flame speed [8], and c) flame-turbulence interaction under supernova conditions. To better study the flame characteristics, the reference frame was chosen such that the flame was stationary in the computational domain. The 3D reacting flow simulations were initialized using the 1D profiles and isotropic turbulence was introduced through the inflow boundary. To understand the effects of various parameters, most notably Da and Ka numbers, as well as the effect of compressibility, several simulations were performed on up to 2048^3 meshes.

There is a complicated phenomenology associated with turbulent flames under type Ia supernova conditions from the suppression of the smallest vortex tubes due to the flame fire polishing, but enhancement of intermediate turbulent scales (Fig. 1), to the rapid acceleration of the flame itself to large velocities, is one of the important open questions related to the supernova modeling. In addition, the fully compressible simulations allowed considering the dynamics of the dilatational motions, neglected in previous studies. These motions are enhanced by the heat addition due to the flame and can cause shock waves that may lead to detona-

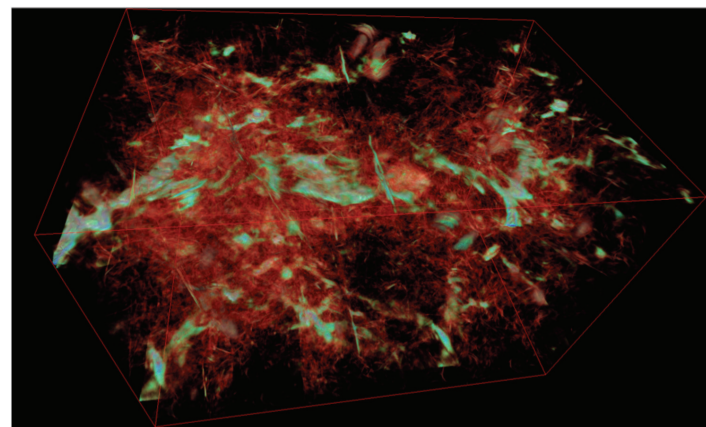


Fig. 2. Dilatational effects, usually neglected in astrophysical calculations, are important for understanding the deflagration-detonation initiation. Our results indicate the presence of shocklets, generated by the turbulence and enhanced by the dilatational motions due to the flame. Sufficiently strong shocklets may be the mechanism for producing detonation in type Ia supernovae.

tion (Fig. 2). Current research focuses on the departures of the turbulence properties from the classical Kolmogorov picture and determining accurate turbulent flame speeds.

**For more information contact
Daniel Livescu at
livescu@lanl.gov.**

- [1] D. Livescu, C.K. Madnia, *Phys. Fluids*, **16**, 2864 (2004).
- [2] D. Livescu, J.R. Ristorcelli, *J. Fluid Mech.* **591**, 43 (2007).
- [3] D. Livescu, J.R. Ristorcelli, *J. Fluid Mech.* **605**, 145 (2008).
- [4] D. Livescu et al., *J. Turb.* **10**, N13 (2009).
- [5] P. Moin, K. Mahesh, *Ann. Rev. Fluid Mech.* **30**, 539 (1998).
- [6] J. Mohd-Yusof, D. Livescu, T. Kelley, *Proc. ParCFD* (2009).
- [7] J. Mohd-Yusof, D. Livescu, T. Kelley, *Proc. CSE* (2009).
- [8] M. Petersen, D. Livescu, *Phys. Fluids* (submitted).

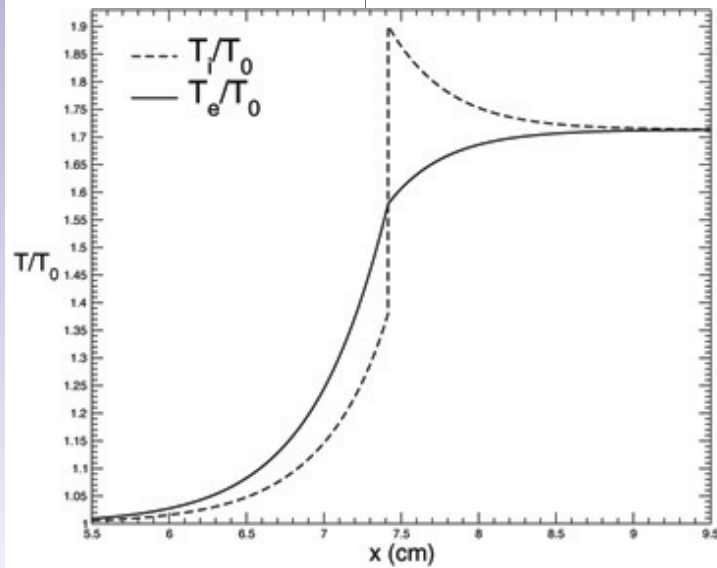
Funding Acknowledgments

LANL Directed Research and Development Program—Exploratory Research (ER)

Shock Wave Structure for an Ionized Plasma

Thomas O. Masser, John G. Wohlbiel, Robert B. Lowrie, CCS-2

Fig. 1. Temperature profiles near Mach 1.7 shock in ionized hydrogen; $R=1.0$.



The importance of electron heat conduction in inertial confinement fusion (ICF) pellets [1,2] and certain astrophysical regimes [3] requires that plasma simulation codes include plasma models with separate ion and electron temperatures. Producing analytic (or semi-analytic) solutions for simple two-temperature plasma models is useful for the verification of the physics algorithms within a simulation code. We study the structure of planar shock waves in a two-temperature model [4-6] of a fully ionized plasma that includes electron heat conduction and energy exchange between electrons and ions. For steady flow in a reference frame moving with the shock, the model reduces to an autonomous system of ordinary differential equations that can be numerically integrated.

The primary focus of this study is to compute and explore the range of possible shock solutions for a model plasma. These solutions may be used to verify hydrodynamic codes that use similar plasma physics models.

We focus only on the interactions of the electrons and ions with a shock moving through a fully ionized gas. We assume that strong Coulomb interactions keep the electrons and ions rigidly coupled, so that the plasma remains electrically neutral. We neglect all radiative effects and treat both the electron diffusivity and electron-ion coupling coefficients as

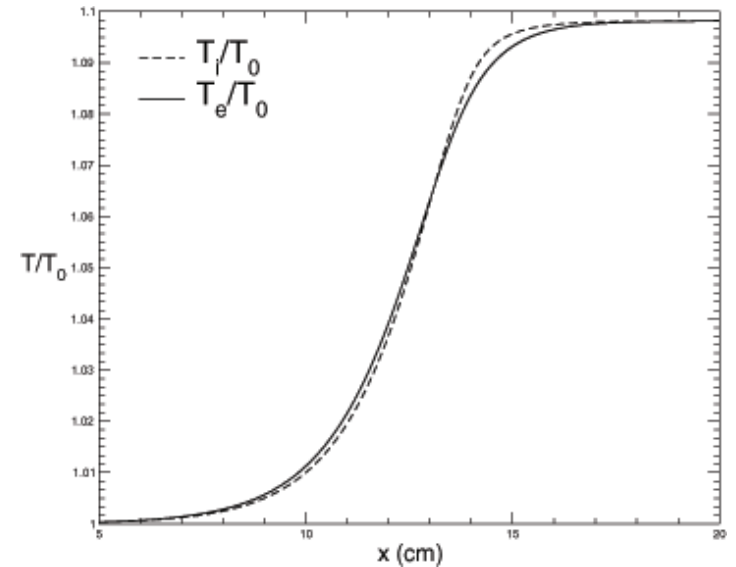


Fig. 2. Temperature profiles near Mach 1.1 shock in ionized hydrogen; $R=1.0$.

constants. While the model will be invalid for determining the true details of plasma shock structures, the solutions described in this study are simple to compute and provide additional insight into the shock structure. The simple model we employ captures the primary effects of shocks on the electron and ion temperatures and may lead to a more complete picture of the range of possible solutions. For instance, we improve on a previous derivation [6] of the boundary between continuous and discontinuous shock profiles. We also show that the ion temperature may continue to increase behind a hydrodynamic shock and achieve a maximum in the region downstream of the shock, similar to an effect seen in radiative shocks [7].

The electron and ion temperature profiles for a shock with Mach number 1.7 passing through ionized hydrogen are given in Fig. 1. The shock differentially heats the ions (protons) and electrons due to the discrepancy in their masses. Electron heat conduction produces precursor heating of the electrons ahead of the shock (left of the shock in Fig. 1). Electron-ion coupling serves to equilibrate the electron and ion temperatures away from the shock.

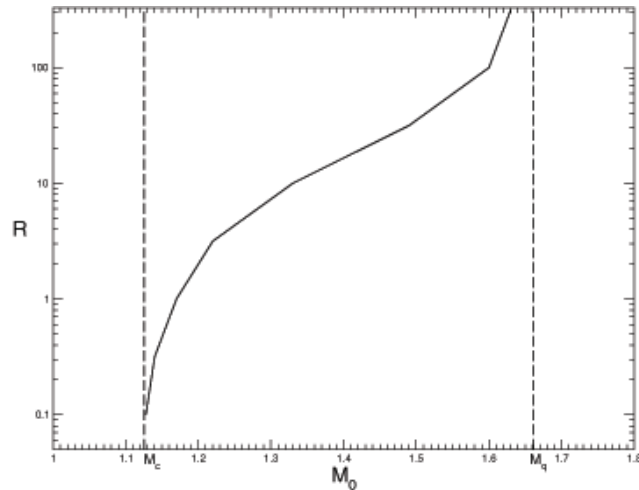


Fig. 3. Boundary between continuous and discontinuous shock solutions.

For sufficiently weak shocks, Imshennik [6] has shown that electron heat conduction will smooth the hydrodynamic discontinuity at the shock so that the temperature profiles will be continuous. In particular, below a critical Mach number (approximately 1.125 for ionized hydrogen), the flow variables (density, pressure, temperatures) will be continuous, regardless of the strength of the electron diffusivity or the electron-ion relaxation. Figure 2 plots the electron and ion temperatures for a shock with Mach number 1.1 passing through ionized hydrogen.

We have shown [8] that two parameters govern the qualitative behavior of these shock solutions: 1) the Mach number of the shock, M_0 , and 2) a term R that is proportional to the product of the electron diffusivity and the electron-ion coupling constant. Figure 3 plots the approximate boundary between continuous and discontinuous solutions for ionized hydrogen in the parameter space (M_0, R). For Mach numbers less than M_c , the shock solutions are always continuous due to the presence of electron heat conduction, as predicted by Imshennik. Above this threshold Mach number, if the electron heat conduction and/or electron-ion coupling are strong enough (i.e., if R is large enough) the shock solutions

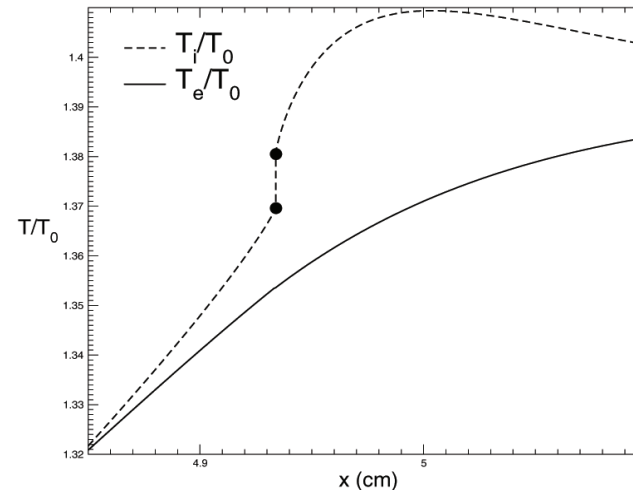


Fig. 4. Temperature profiles near Mach 1.4 shock in ionized hydrogen; $R=10.0$.

remain continuous; otherwise, the solutions exhibit an embedded hydrodynamic shock, as in Fig. 1. Our numerical results indicate that a second threshold Mach number exists; for shocks with Mach number greater than M_q , the solutions are always discontinuous, i.e., the solutions exhibit an embedded hydrodynamic shock regardless of the strength of the electron heat conduction or electron-ion coupling.

Physically, the presence of electron heat conduction and electron-ion coupling leads to regions of continuous compression of the plasma ahead of and behind the shock. When the diffusion and relaxation effects are weak relative to the hydrodynamic shock, most of the plasma compression occurs at the embedded hydrodynamic shock. When these effects are strong relative to the hydrodynamic shock, significant continuous compression occurs in both the precursor and relaxation regions and not at

the shock. A maximum ion temperature may occur away from the shock when the energy flowing into the ions from the continuous compression is balanced by the energy flowing from the ions into the electrons. Figure 4 exhibits a shock solution with a postshock maximum in the ion temperature.

**For more information contact
Thomas Masser at
tmasser@lanl.gov.**

- [1] S. Atzeni, J. Meyer-Ter-Vehn, *The Physics of Inertial Fusion*, Oxford University Press (2004).
- [2] J. Lindl, *Inertial Confinement Fusion*, Springer Verlag (1998).
- [3] R. Drake, *High Energy Density Physics*, Springer Verlag (2006).
- [4] Y. Zel'dovich, *Soviet Phys. JETP* **5**, 919 (1957).
- [5] V. Shafranov, *Soviet Phys. JETP* **5**, 1183 (1957).
- [6] V. Imshennik, *Soviet Phys. JETP* **15**, 167 (1962).
- [7] R. Lowrie et al., *Shock Waves* **18**, 129 (2008).
- [8] T. Masser et al., in preparation (2010).

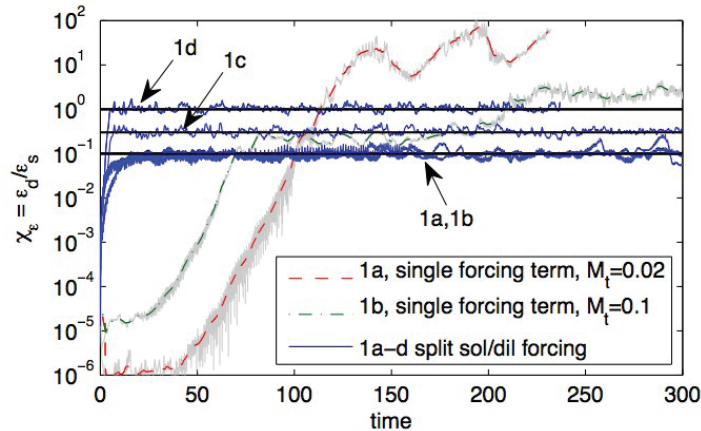
Funding Acknowledgments

DOE, NNSA Advanced Simulation and Computing (ASC) Program

A Realistic and Stable Method of Forcing for Compressible Isotropic Turbulence

Mark R. Petersen, Daniel Livescu, Sumner Dean, Jamaludin Mohd-Yusof, CCS-2

Fig. 1. Evolution of the dilatational-to-solenoidal dissipation ratio, for single term and the new split forcing method. The single-term method is unstable, and dilatational energy can grow without bound. Four simulations with our proposed split forcing method show that statistics quickly adjust to the imposed value (horizontal black lines).



Numerical simulations to investigate the behavior of isotropic turbulence play an important role in turbulence research, especially since the advances in supercomputing power now allow simulations at Reynolds numbers comparable with or even larger than those obtained in laboratory experiments. If a forcing term is added to the equations, then the flow could reach a statistically stationary state, where the injection rate (usually at large scales) is equal to the rate of energy dissipated at small scales. Forcing has several advantages—the Reynolds number of the simulation can be increased considerably, statistics can be averaged over time, decreasing statistical variability due to transient effects, and natural systems are usually forced.

Forcing in nature is due to large-scale effects, for example, solar-induced, buoyancy-driven convection in the atmosphere. The characteristics of turbulence at much smaller scales is thought to be independent of the nature of the forcing, which is why turbulence is often studied in an idealized, triply periodic domain.

We have introduced a method of forcing for compressible, isotropic turbulence where one may precisely control the stationary-state values of the dissipation (and thus the Kolmogorov scale) and the ratio of dilatational (compressive) to solenoidal (rotational) energy [1]. This method is based on a

linear-forcing method, where the forcing term is simply proportional to the momentum, that has been successfully implemented for incompressible flows [2,3].

Linear forcing is physically realistic because it is similar to the natural Reynolds shear stress production mechanism in the turbulent kinetic energy equation. The single-term linear forcing used in current incompressible simulations proved to be unstable for compressible flows because there is no control over the dilatational energy, which grows without bound (Fig. 1). In order to overcome this issue, we developed a two-term forcing method that controls the solenoidal and dilatational dissipation separately.

Forcing may be applied only at low wavenumbers (large scales) or full spectrum (all scales). A comparison of these methods in the context of linearly forced compressible flows showed that low-wavenumber forcing achieves Reynolds numbers that are nearly twice as high, leading to increased range of dynamically relevant scales at the same resolution. The low-wavenumber method also allows us to avoid the computational cost of a full Fourier transform, as one may simply sum the low-wavenumber modes in spectral space.

We research the nature of compressible turbulence using direct numerical simulation (DNS), where all scales down to the viscous

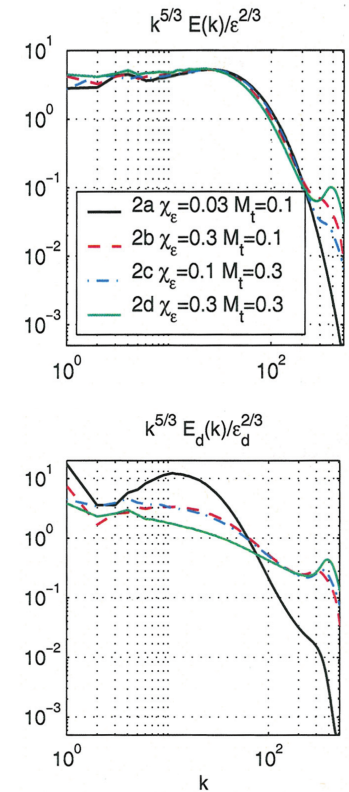


Fig. 2. Compensated energy spectra from four 1024^3 isotropic simulations: total kinetic energy (top), and its dilatational component (bottom).

dissipation range are fully resolved [4]. The two-term linear, low-wave forcing method has allowed us to conduct a large study of $M_t - \chi$ parameter space, where M_t is the turbulent Mach number and χ is the ratio of dilatational to solenoidal kinetic energy. These parameters exhibit the effects of compressibility; low M_p , low χ simulations are nearly incompressible, while high M_p , high χ simulations have strong density gradients and shock waves. Long-time averages of statistics in $M_t - \chi$ parameter space show how these quantities vary. This type of data is useful for testing and developing models for the effects of compressibility in engineering-scale simulations where the smallest scales are not resolved.

For this study, four simulations with 1024^3 grid points were conducted on 1024 processors on Purple, a supercomputer at LLNL. These are the highest resolution compressible Navier-Stokes DNS to date, and achieve Taylor Reynolds numbers greater than 300. In addition, eighteen simulations with Taylor Reynolds numbers of ~ 100 were conducted on the San Diego Supercomputer Center's Blue Gene system. These simulations cover a range of $M_t = 0.02$ -0.3 and $\chi = 0.0$ -1.0. Results show that the total energy spectrum is close to the Kolmogorov 1941 $k^{-5/3}$ law in the inertial range for all parameters, while the dilatational (compressive) spectrum is steeper than $k^{-5/3}$ (Fig. 2), but varies based on the parameters.

There are many differences between compressible turbulence and the more commonly studied incompressible turbulence. In many applications, such as ocean dynamics and large-scale atmospheric motion, the velocity fluctuations are small compared with the speed of sound, the density is nearly constant, and the irrotational effects are small. However, as the velocity fluctuations become comparable to the sound speed and/or exothermic reactions take place, new phenomena occur—shock waves, localized expansion and contractions providing a distinct dissipation mechanism, new nonlinearities leading to additional scale coupling, and strong fluctuations of the thermodynamic quantities. The characteristics of compressibility can be observed in visualizations of divergence and vorticity, where large variability and gradients can be seen (Figs. 3 and 4). Shocklets are short-lived, isolated shock waves that

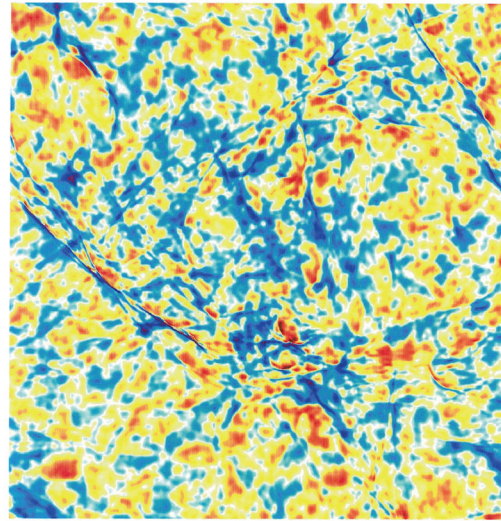


Fig. 3. Cross section of the divergence in a 1024^3 compressible isotropic simulation. Purple lines show shocks in the flow.

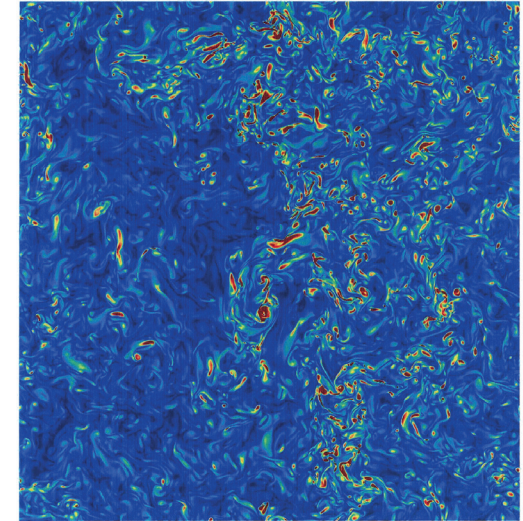


Fig. 4. Cross section of vorticity magnitude show an active eddy field.

occur in these areas, and are thought to be responsible for many of the effects of compressible turbulence [5].

For more information contact Mark Petersen at mpetersen@lanl.gov.

-
- [1] M.R. Petersen, D. Livescu, "Forcing for statistically stationary compressible isotropic turbulence," *Phys. Fluids*, (in review).
 - [2] T.S. Lundgren, *VL Annual Research Briefs*, **461**, (2003).
 - [3] C. Rosales, C. Meneveau, *Phys. Fluids* **17**, 095106 (2005).
 - [4] D. Livescu et al., *J. Fluid Mech.* **450**, 35 (2002).
 - [5] R. Samtaney et al., *Phys. Fluids* **13**, 1415 (2001).

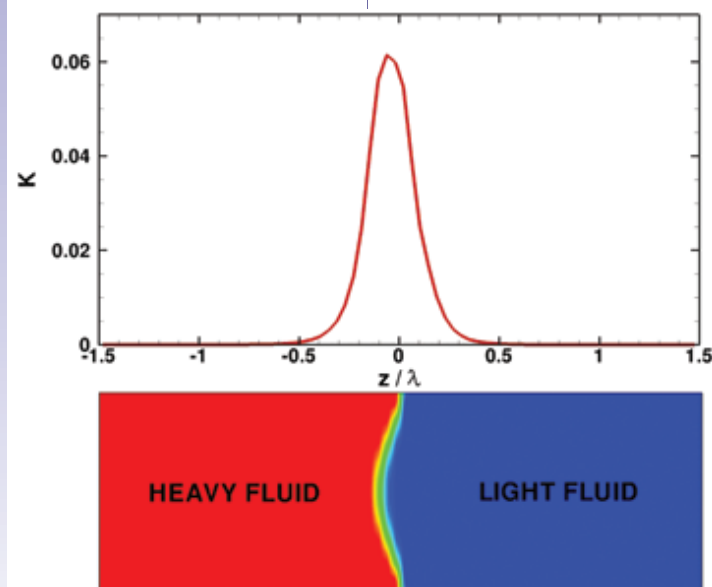
Funding Acknowledgments

DOE, NNSA Advanced Simulation and Computing (ASC) Program

On the Early-time Kinetic Energy Profile for the Rayleigh-Taylor Instability

Bertrand Rollin, Malcolm J. Andrews, CCS-2

Fig. 1. Top: kinetic energy profile soon after the beginning of the nonlinear growth. K : kinetic energy; λ : wavelength of the initial perturbation. Bottom: interface profile. $A_T=0.5$.



The Rayleigh-Taylor instability is an interfacial fluid instability that leads to turbulence and turbulent mixing. It occurs when a light fluid is accelerated into a heavy fluid [1,2] because of misalignment between density and pressure gradients. The Rayleigh-Taylor instability plays a key role in a wide variety of contexts such as supernovae explosions, salt dome formation, and oceanic flows, as well as in technological applications such as rotating machinery or in the implosion phase of inertial confinement fusion (ICF).

Traditional research in turbulence assumes that turbulent flows have no memory of the initial conditions, and turbulence develops to a universal self-similar state. However, recent research has shown that initial conditions are of great importance during the growth of the turbulent Rayleigh-Taylor instability [3,4]. As a result, the need for accurate initial conditions in turbulence models has now become clear.

Kinetic energy is a fundamental parameter in most common turbulence models. In this work, our objective was to study the early-time evolution of the kinetic energy profile in the case of a single-mode initial perturbation. It is indeed important to have a good description of this basic

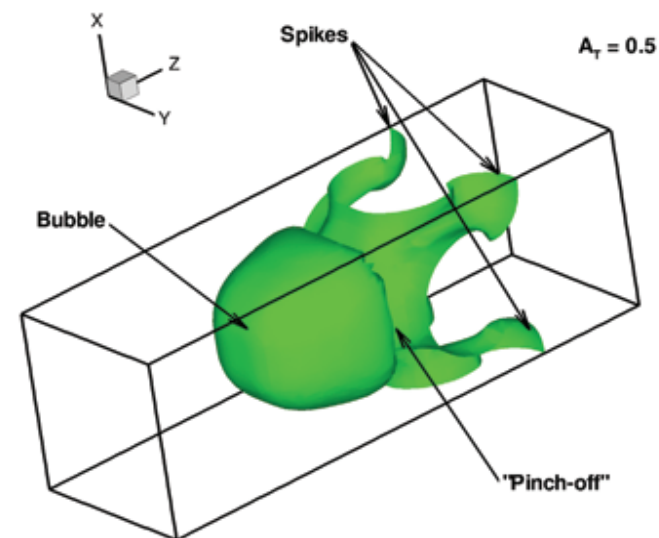


Fig. 2. Interface between the heavy and light fluid at $A_T=0.5$.

case, since most of the complex models describing the multimode turbulent instability are built on a single-mode model (c.f., Fourier decomposition).

Simulations of a single-mode initial perturbation at low, intermediate, and large Atwood numbers (density contrast) were performed using an incompressible MILES code [5], RTI3D. During the linear and early nonlinear stage of the instability, the kinetic energy profile is simple as the flow remains relatively symmetrical (Fig. 1), and reflects the sinusoidal nature of the perturbation. We observe a peak of kinetic energy located around the centerline of the initial sine perturbation. The profile prediction during the linear (actually exponential growth phase) is then straightforward using velocity predictions from theoretical models. Figure 2 shows a typical shape of the interface (computed using RTI3D) between the heavy and the light fluid after some time in the nonlinear regime ($h_B/\lambda \sim 0.6$, where h_B is the height of the bubble). The bubble (blob of light rising into the heavy fluid) is now well formed. The vortices on the spikes (blobs of heavy fluid falling into the light fluid) side are well developed, and

we observe a pinch-off region at about the centerline. At this stage of the nonlinear regime, theoretical models are still able to provide a good estimate of the position of the tips of the bubbles and the spikes at low and intermediate Atwood numbers [6]. However, there are to the authors' knowledge no models predicting the velocity distribution in the mixing region, and hence no model to predict the kinetic energy profile. The kinetic energy profile from RTI3D is displayed on the bottom pictures of Fig. 3 for different Atwood numbers, and the shape of the instability colored by the instantaneous kinetic energy is displayed on the top pictures. All profiles show two local maxima corresponding to the vortices located at the pinch-off and near the tip of the spikes. However, if the instantaneous kinetic energy is higher near the spikes (top row Fig. 3), the azimuthally averaged kinetic energy (bottom row Fig. 3) displays a higher value in the pinch-off region, which occupies a larger volume. As the Atwood number increases, the asymmetry between spikes and bubbles increases. For larger Atwood number, the intensity of the vortices is more inhibited by the density contrast, resulting in a smaller value of kinetic energy. Figure 4 shows that the local maxima in the kinetic energy profile are mainly due to the contribution of the vortex ring of the pinch-off region and the smaller vortex rings located near the tip of the spikes. The observations made in this work will provide us with guidance for deriving a model for kinetic energy profile prediction at early-time for the turbulent Rayleigh-Taylor instability.

For more information contact Bertrand Rollin at
bertrand@lanl.gov.

- [1] Lord J. Rayleigh, *Proc. Lond. Math. Soc.* **14**, 170 (1883).
- [2] G.I. Taylor, *Proc. Royal. Soc. Lond. A* **201**, 192 (1950).
- [3] P. Ramaprabhu, G. Dimonte, M.J. Andrews, *J. Fluid Mech.* **536**, 285 (2005).
- [4] A. Banerjee, M.J. Andrews, *Int. J. Heat Mass Transfer* **52**, 3906 (2009).
- [5] M.J. Andrews, *Int. J. Numer. Meth. Fluids* **21**, 205 (1995).
- [6] V.N. Goncharov, *Phys. Rev. Lett.* **88**, 134502 (2002).

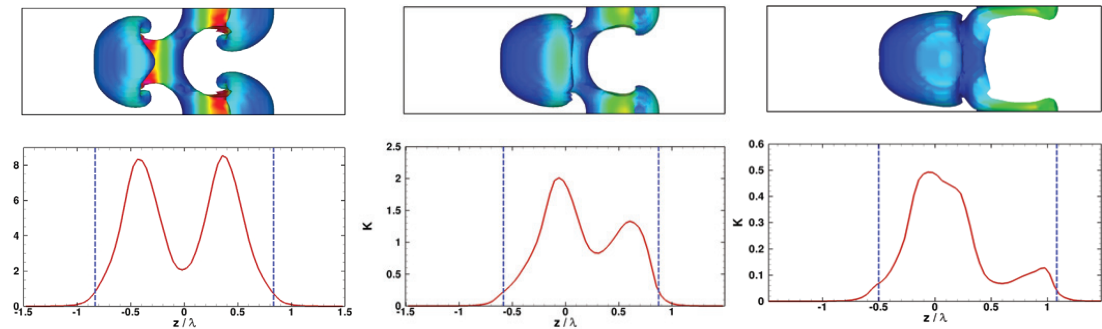


Fig. 3. Top: View in the xz -plane of the interface between the heavy and the light fluid. **Bottom:** kinetic energy profile (red). The dashed blue lines indicate the position of the tips of the bubble and spikes. $A_T=0.0$ (left); $A_T=0.5$ (middle); $A_T=0.9$ (right).

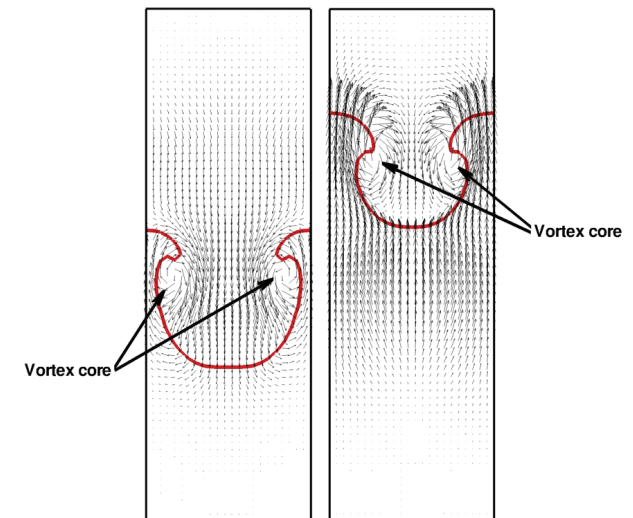


Fig. 4. Slice of the velocity distribution in the mid-plane $y = \lambda/2$ (left) and the plane $y=0$ (right). The interface is represented by the red line.

Funding Acknowledgments

LANL Directed Research and Development Program

Improved Interface Treatment in Arbitrary Lagrangian-Eulerian Simulations using the Moment-of-Fluid Method

M. Kucharik, FNSPE, Czech Technical University in Prague; Rao V. Garimella, Samuel P. Schofield, Mikhail J. Shashkov, T-5

We have successfully extended the Moment-of-fluid (MOF) interface reconstruction and tracking method to Arbitrary Lagrangian-Eulerian (ALE) compressible, multimaterial, multiple equation-of-state hydrodynamic simulations. In these simulations, the MOF method shows significant improvement over existing interface reconstruction methods.

In Lagrangian hydrodynamic simulations the equations of conservation of mass, momentum, and energy are solved in a Lagrangian reference frame where the computational mesh moves with the material. However, for many flows, vorticity or shear in the flow may ultimately tangle the mesh and terminate the simulation. This has led to the development of ALE methods where the mesh may move with a different velocity than the computed flow field in order to preserve mesh quality.

ALE simulations consist of three steps: 1) a Lagrangian phase, where the equations are solved in the Lagrangian reference frame, 2) a mesh smoothing phase, where a better quality mesh is determined based on the initial mesh or the mesh at the end of the Lagrangian step (this step may be done after a set number of Lagrangian steps or when a mesh quality threshold is reached), and 3) a remap phase, where conserved variables are interpolated from the Lagrangian mesh of step 1 to the smoothed mesh from step 2 in a conservative manner.

In a pure Lagrangian calculation material interfaces are preserved through the calculation assuming they are originally at cell boundaries. However, for ALE simulations interfaces will no longer be aligned with the mesh and an interface treatment is necessary. ALE simulations have traditionally tracked the fractional volumes and used

interface reconstruction [1] to approximate the interface positions in each mesh cell when required. In the interface reconstruction process, a line segment representing the interface within the cell is positioned to match the prescribed material volume. Interface reconstruction methods are primarily differentiated by how the orientation of the line is determined. The widely used first-order-accurate Youngs method [2] uses the negative gradient of the volume fractions as the outward normal. It works well for two materials but has an artificial material order dependency when reconstructing interfaces for more than two materials. The MOF method [3], developed at LANL, utilizes the first moments in addition to fractional volumes of the materials to reconstruct the interface. It is material-order independent and can correctly locate each material in a cell providing true sub-cell resolution.

However, in compressible flow simulations the fractional volumes are not necessarily a Lagrangian invariant and the time evolution of the first moments must account for the volumetric change. This has impeded the adoption of MOF in compressible and ALE simulations, although MOF has shown significant improvements in incompressible flow simulations [4].

In this work, we have developed a method to update the centroid in the Lagrangian phase that is provably second-order accurate, and we have developed a method to update the material moments in the remap phase that preserves the accuracy of the MOF method.

To update material centroids during the Lagrangian phase, the material centroid in a cell is assumed to be constant in barycentric coordinates before and after the Lagrangian step. We have shown [5] this provides a second-order accurate approximation to the true evolution of a material centroid. This approach is very general, working for arbitrary convex polygonal mesh cells.

During the remap step, the material moments on the new, smoothed mesh are calculated using an intersection-based remap written in flux form to provide conservation to machine accuracy. The overall method has been shown to be robust and accurate.

In the triple-point problem, a high-pressure region in the left of the 7×3 domain expands into two layers of different gases as shown in Fig. 1. Here ρ is the density, p is the pressure, \mathbf{u} is the velocity and γ is the adiabatic exponent. The top, blue material compresses more readily, allowing the shock wave to travel faster than in the bottom green material. As a result, strong rotation is generated at the interface between the three materials. This vortex creates problems for the Lagrangian simulation, but the ALE simulation handles it gracefully – but it must reconstruct the interfaces.

Because Youngs method is material-order dependent, that is, the reconstructed interfaces depend on the order they were processed, different material orderings may give different results. This is shown in the zoomed-in portion of Fig. 1 where different orderings have caused mixing or separation of the material interfaces. MOF not only accurately preserves the interfaces, it can evaluate all possible material orderings and choose the optimal configuration locally.

For more information contact Samuel P. Schofield at sams@lanl.gov.

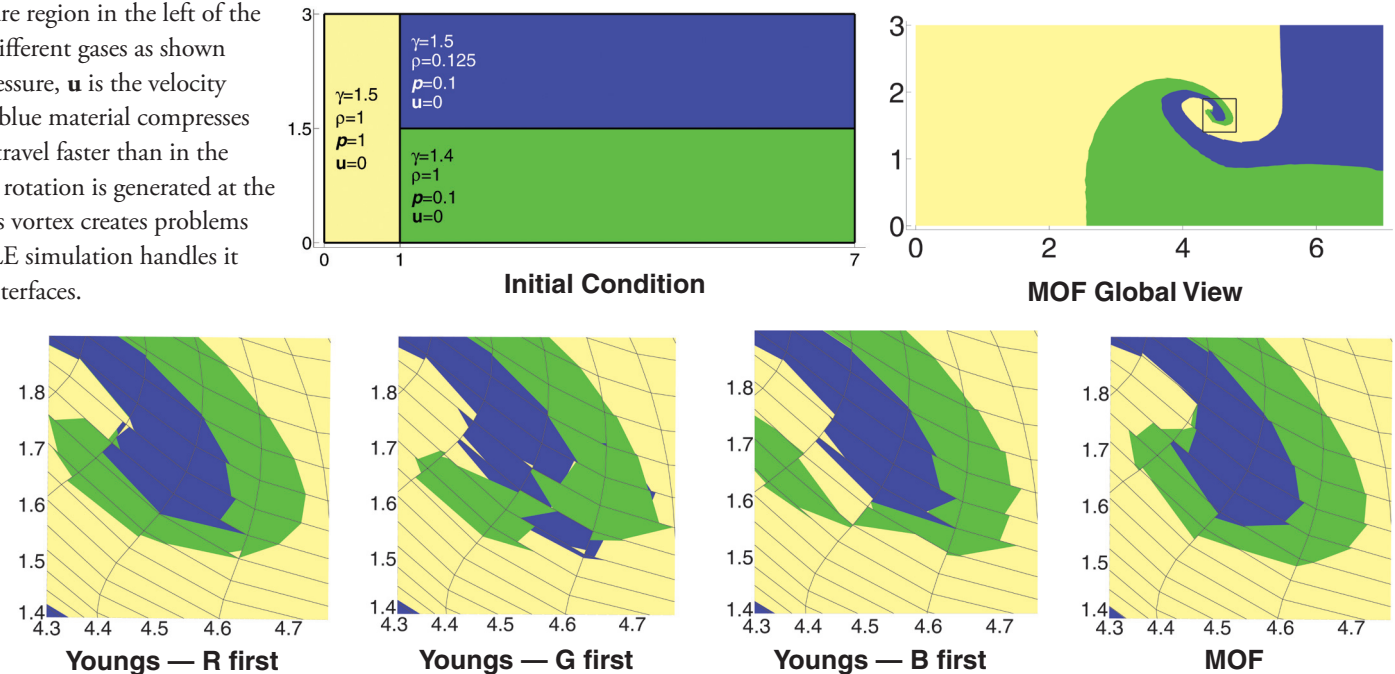


Fig. 1. The initial condition and material positions for the triple-point problem at time $T=5.0$. The mesh was relaxed using the Winslow method [6] after every 20 Lagrangian steps. The zoomed-in portion shows the material positions with different material orderings and with the MOF method. The mesh consists of 140×60 cells with material interfaces initially aligned with the mesh.

- [1] W.J. Rider, D.B. Kothe, *J. Comput. Phys.* **141**, 112 (1998).
- [2] D.L. Youngs, *Numerical Methods for Fluid Dynamics*, 273 (1982).
- [3] V.D. Dyadechko, M. Shashkov, *J. Comput. Phys.* **227**, 5361 (2008).
- [4] S.P. Schofield et al., *Int. J. Numer. Meth. Fluids*, DOI: 10.1002/fld.2108 (2009)
- [5] M. Kucharik et al., *J. Comput. Phys.* **229**, 2432 (2010).
- [6] A.M. Winslow, Technical Report UCRL-7312, LLNL (1963).

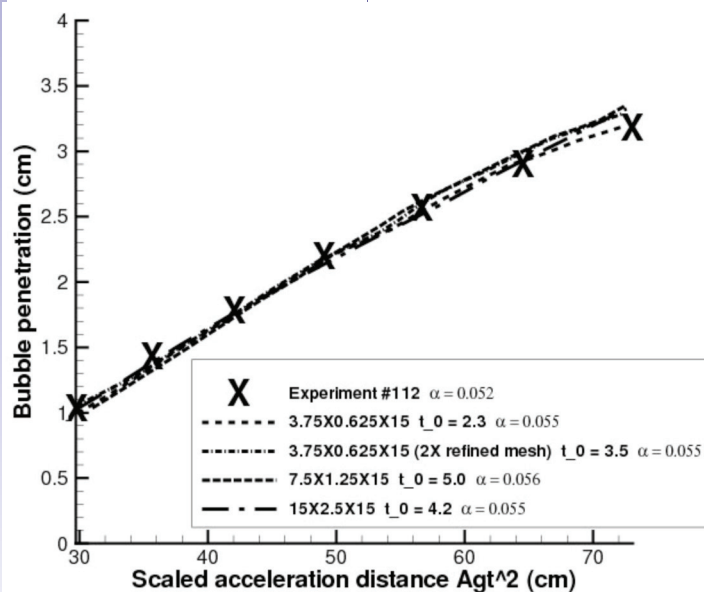
Funding Acknowledgments

LANL Directed Research and Development Program

Turbulent Mixing Simulations with Verification and Validation

James Glimm, David H. Sharp, ADTSC

Fig. 1. Smeeton-Youngs experiment #112 (salt-fresh water) compared with several simulations with different mesh and statistical resolutions.



Turbulent mixing is central to a number of LANL problems, including inertial confinement fusion (ICF) and turbulent combustion. Classic acceleration-driven instabilities (Rayleigh-Taylor for continuous acceleration and Richtmyer-Meshkov for shock acceleration) have been a numerical challenge for 50 years. A satisfactory solution should use a compressible code, satisfy verification and validation (V&V) standards, achieve mesh-converged probability distributions for temperature and mixture concentrations, not require full resolution direct numerical simulation (DNS) of all turbulent length scales, apply over a range of physical parameters, including high Schmidt numbers, and provide sufficient intellectual understanding that others can follow and improve.

Our progress report, based on recent results from a LANL-Stony Brook University collaboration, outlines results that achieve these six goals.

The compressible code is the LANL-Stony Brook code FronTier, whose main two innovative features are front tracking, to eliminate numerical mass diffusion across fluid interfaces, and dynamic subgrid scale (SGS) models, to represent turbulent fluctuations at levels below the grid resolution.

The SGS models are free of adjustable parameters, and the entire simulation is thus free of adjustable parameters. Such SGS models are standard, but our use of them in flows with flow gradients that are strong at a mesh level appears to be original.

The main results are contained in a series of recent papers [1-4]. The summary here will emphasize the sixth point—the intellectual and conceptual innovations needed (going beyond the capabilities of an advanced computer code) to obtain these results.

Commonly, the Rayleigh-Taylor mixing simulations show a factor of two or more discrepancy between simulation and experiment. Conventionally, this discrepancy is attributed to long wavelength noise in the experiments, but not measured—hence the simulations and (most) experiments should not be compared. The Mueschke-Andrews experiments (hot-cold water flowing over a splitter plate) are an exception—measured initial data shows very substantial long wavelength initial perturbations. We duplicated the Mueschke-Schilling fully resolved DNS simulation results, with agreement between experiment and simulation. We extended this agreement to the high Schmidt number case (salt-fresh water over a splitter plate), using a fully compressible code with SGS models, and a large eddy simulation (LES), i.e., not requiring full DNS resolution. The high Schmidt number case is inaccessible to the above DNS methods.

However, not all experiments have such a high level of noise. Our analysis of the Smeeton-Youngs rocket rig (fresh-salt water) experiments suggested that they were relatively free of long wavelength noise. We simulated these, in agreement with experiment, and for the first time, with a simulation going through all experimental data points (see Fig. 1).

Our results show a clear dependence of the mixing rates on the Schmidt number. A common simulation strategy is known as Implicit Large Eddy Simulation (ILES), whereby the numerical algorithm supplies what it will for the subgrid effects, while the physical transport parameters are either zero or so small as to leave the transport effects underresolved on any feasible grid level, and thus negligible relative to numerical transport. The Schmidt number

is the ratio of viscosity to mass diffusivity. If both numerator and denominator are set to zero, the ratio is the indeterminate expression 0/0. If both are set by numerical algorithms and code artifacts, then it is a code-dependent numerical Schmidt number that is determining the answer. The result (and this was observed in practice) is two well-documented codes converging numerically to two distinct answers (see [5-7] for this and related results). Clearly this is a major challenge for V&V. The SGS models and a valid LES simulation cure this ambiguity, and allow mesh convergence, even for such sensitive and important observables as the joint probability distributions for temperature and species concentration, or for a chemical reaction rate.

The mathematical literature contains non-uniqueness theorems for the compressible Euler equations, i.e., fluid equations with fluid transport set to zero. Thus V&V for Implicit Large Eddy Simulations (ILES) has hit a mathematical no-go theorem. Obviously, a change in strategy is needed, and this is what we provide.

In comparing the Smeeton-Youngs to the Mueschke-Andrews fresh-salt water Rayleigh-Taylor mixing experiments, we found six significant differences in the experimental conditions, and all of them had a significant effect on the mixing rate (α). Beyond the fluid transport parameters and long-wave noise (in the Mueschke-Andrews experiment only), we found significant influence from the change in Grashov number, initial mass diffusion layer width, and experimentally induced vs. ideal short wave length perturbations.

Thus a central conclusion of our study is that there is nothing universal about Rayleigh-Taylor mixing. Many details of the experiment are important. When modeled correctly, nearly perfect agreement between simulation and experiment can be obtained. Combined with mesh convergence, this is V&V.

We also propose a scientific program for recovery in an ILES world. If front tracking to control or limit numerical mass diffusion is not part of a game plan, then it is necessary to increase the numerical viscosity so that the ratio of numerical viscosity to numerical mass diffusion agrees with the physical ratio, i.e., the numerical Schmidt number must equal the physical Schmidt number for a successful

Table 1. Mesh and model errors for the joint PDFs for temperature and species concentrations and for typical gas and liquid parameter values for fluid transport. Comparison of coarse and medium to fine meshes (c to f and m to f) using sup norm for the difference of the probability distribution functions and an L1 norm for the comparison of model to simulation.

	liquid			gas		
Re	c to f	m to f	model	c to f	m to f	model
≈ 300	0.24	0.13	0.06	0.08	0.06	0.06
≈ 6000	0.07	0.06	0.07	0.09	0.05	0.05
$\approx 600K$	0.26	0.06	0.07	0.22	0.09	0.08

ILES simulation. The implementation of this plan, which is not without difficulties, will require a calibration step for the selection of the value for numerical viscosity. Our recommended route is to abandon this narrow interpretation of ILES, and to combine the multiphysics capabilities of ILES codes with the high quality geometrical resolution of steep concentration gradients in FronTier. To this end, a FronTier API, an application programming interface designed to support code merger is under development.

Going beyond the study of α , i.e., the outer envelope of the mixing region, we have attained mesh convergence for probability density functions (PDF) of the temperature and species concentration [1]. Due to the need for higher levels of mesh refinement for this convergence study and the plan to study the variation of physical parameters, we reduced the computational burden with a 2D study, which was of a circularly imploding Richtmyer-Meshkov instability, studied past the time of reshock. With the LES SGS regularization the convergence issues were tractable, and the

mesh convergence of the PDFs was obtained. Moreover, a simple theoretical model, based on the geometry of the 50% concentration isosurface (as computed) and the laminar and turbulent (as computed) diffusion coefficients was used as an independent check on the mixing PDFs, and gave errors comparable to the mesh convergence mesh errors (see Table 1).

For more information contact David H. Sharp at dcso@lanl.gov.

- [1] H. Lim et al., *Phys Fluids*, (submitted for publication); LA-UR 08-07725 (2008).
- [2] H. Lim et al., *J. High Energy Physics*, (in press); LA-UR 09-01364 (2009).
- [3] H. Lim et al., *Proc. Natl. Acad. Sci. (USA)* (submitted for publication); LA-UR 09-06333 (2009).
- [4] H. Lim et al., *Physica Scripta*, (submitted for publication); LA-UR 09-07240 (2009).
- [5] T. Masser. *Breaking Temperature Equilibrium in Mixed Cell Hydrodynamics*. Ph.D. Thesis, State University of New York (2007).
- [6] H. Lim et al., *Acta Mathematica Applicatae Sinica* **24**, 355 (2008).
- [7] H. Lim et al., *Methods Appl. Mech. Eng.* 197, 3435 (2008).

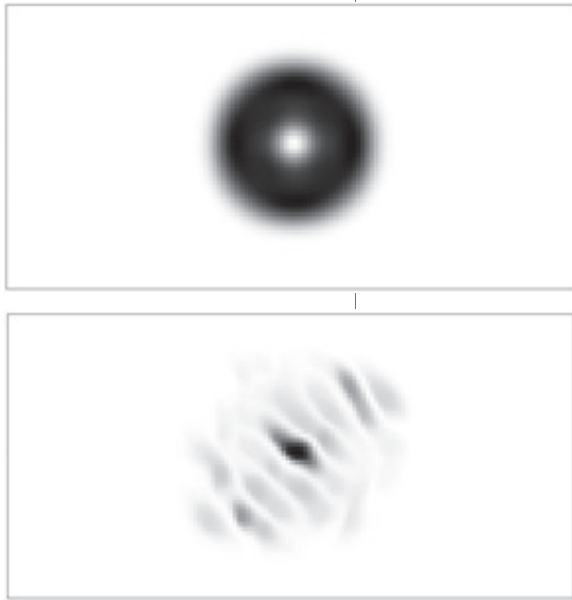
Funding Acknowledgments

- DOE, NNSA Predictive Science Academic Alliance Program (PSAAP) through Stanford University.
- NNSA grant
- Army Research Office (ARO)

High-order Divergence-free MHD Solver for Turbulence Simulation

Shengtai Li, T-5

Fig. 1. The divergence-free method (top) preserves the magnetic field-loop better than the nondivergence-free method (bottom).



Turbulence is a common phenomenon in plasmas. It is convenient to treat turbulent plasmas in a magneto-hydrodynamic (MHD) framework. MHD studies the dynamics of electrically conducting fluids. Magnetic fields have a nice property called divergence-free condition, which means that the divergence of the magnetic fields is always zero, for conservation of the magnetic flux.

The MHD turbulence simulations demand methods with both high-order accuracy and a divergence-free preserving property. The order of accuracy, which is also called the rate of convergence of simulation results to the true solutions, is one of main factors of the efficiency of the method. The low-order method demands a much finer grid resolution than the high-order method in order to achieve the same accuracy. Numerical experiments show that the eighth-order scheme can save CPU time by a factor of 85, and data storage by a factor of 64, over the second-order scheme in achieving a comparable effect in turbulence simulations. The divergence-free preserving method is also essential in order for MHD simulations to produce the correct dynamics (see Fig. 1). Due to the lack of highly accurate simulation software, the research and development on compressible MHD turbulence is still in its infancy. We are aware of no higher than a second-order divergence-free method thus far in MHD simulations.

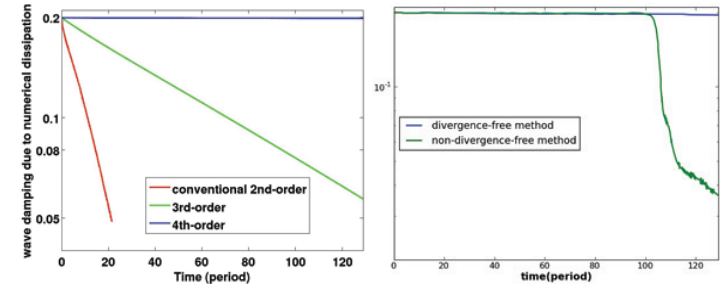


Fig. 2. The left plot shows the fourth-order method has much lower numerical dissipation than the third- or traditional second-order methods. The right plot shows the nondivergence-free method can introduce large dissipation at certain times during simulations.

The traditional numerical simulations use only one grid to cover the whole physical domain. The magnetic field at the cell boundaries can numerically become discontinuous due to the numerical interpolation. The unphysical discontinuity of the magnetic field can trigger numerical instability [1] and crash the simulation.

Inspired by the work in [2], we propose a new methodology to develop numerical methods for MHD simulations. Rather than using only one grid, we generate an overlapping dual grid from the primal grid. By solving the MHD equations on both the primal and dual grids, the cell boundaries of one grid fall naturally within the dual grid, the magnetic fields become continuous in the dual grid, and the flux is easy to calculate by using the dual grid information. The accuracy order of the method is equal to the order of the reconstruction on a cell, which can be arbitrarily-high order accuracy. The divergence-free condition can also be easily achieved by developing high-order, divergence-free reconstruction over a cell.

We first developed a third-order divergence-free method for ideal MHD in 2D [3]. With the help of a dual grid, we get rid of the cell-centered magnetic field and related spatial averaging, which is needed by conventional method. We can compute the magnetic field with arbitrarily high-order accuracy. We have constructed a compact third-order, divergence-free scheme for MHD simulations. To our knowledge, this is the first verified higher than second-order scheme

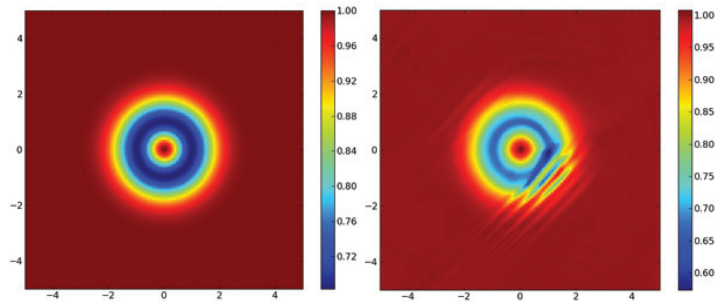


Fig. 3. The divergence-free method preserves the vortex very well, whereas the nondivergence-free method introduces a large error in the vortex.

with the divergence-free property for MHD simulations. For MHD flows that contain shock and contact discontinuity, we proposed an essentially nonoscillatory reconstruction to preserve the monotonicity of the fluid variables. This limited reconstruction uses only adjacent nearest neighbors via a hierarchical procedure.

We have also developed a 4th-order method using the same framework as the 3rd-order method [4]. The 4th-order, divergence-free reconstruction requires a constraint in addition to the dual overlapping grid. Numerical verification shows that our method achieves 4th-order accuracy and is divergence-free for the magnetic fields. We have found that the 4th-order method is roughly 32 times more efficient than the 3rd-order method in order to achieve the same accuracy as for 2D problems. Figure 2 shows the 4th-order method has much lower numerical dissipation than the 3rd-order method and conventional 2nd-order method.

We have found that the divergence-free property is crucial in our high-order method, especially when limited reconstruction is used. The right plot of Fig. 2 shows that without the divergence-free condition, the wave is suddenly damped due to the large divergence error in the numerical method. Figure 3 shows the pressure for a vortex advection in a periodic domain. After five periods, the vortex solved by the divergence-free method is preserved very well, whereas the nondivergence-free method has a large error produced.

Figure 4 shows that three linear waves propagate and interact with each other before they become turbulence. The nondivergence-free solution has a lot of numerical noise in the wave profile, which is introduced by a large divergence error in the magnetic field.

We have extended our 2D 3rd- and 4th-order methods to 3D. The extension is not trivial and will be documented in our future paper. Numerical verifications show that we have achieved the expected order of accuracy for both the divergence-free and nondivergence-free methods.

Comparing the efficiency and accuracy of the 3rd- and 4th-order methods for 3D, we have found that the 4th-order method is 64 times more efficient than the 3rd-order method in order to achieve the same accuracy for 3D problems.

For more information contact Shengtai Li at sli@lanl.gov.

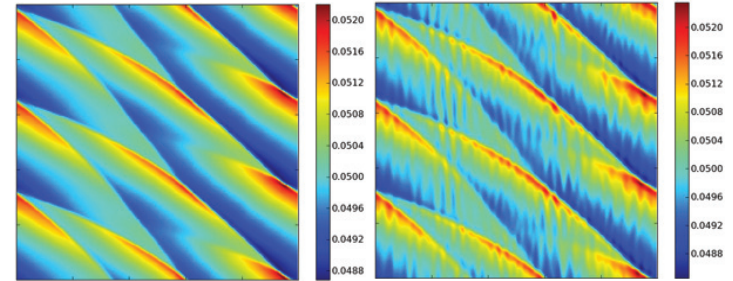


Fig. 4. The nondivergence-free method (right) introduces much noise to the solutions in the linear waves interaction and propagation problem, whereas the divergence-free method (left) retains a clean wave profile.

- [1] T.J. Barth, *IMA volume on Compatible Discretization*, Springer-Verlag (2005).
- [2] Y. Liu et al., *Commun. Comput. Phys.* **2**, 933 (2007).
- [3] S. Li, *JCP* **227**, 7368 (2008).
- [4] S. Li, "The fourth-order divergence-free method for magneto-hydrodynamic flows," (to be submitted to *JCP*)

Funding Acknowledgments

LANL Directed Research and Development Program

A New Monotone Finite Volume Method for Advection-Diffusion Problems

Konstantin Lipnikov, Daniil Svyatskiy, T-5; Yuri Vassilevki, Institute of Numerical Mathematics of the Russian Academy of Sciences, Russia

The advection and diffusion processes are among the main ingredients of the modeling of flow and transport in a porous media. These processes are tightly coupled, so in this work we consider the steady-state advection-diffusion problem. The maximum principle and local mass conservation are the essential properties of this problem. In spite of the simple structure of the underlying equation, it is very difficult to incorporate these properties into a numerical scheme. Highly anisotropic media properties, distorted meshes, and advection-dominated flows may cause nonphysical oscillations of numerical solutions in the vicinity of internal shocks and boundary layers. These oscillations may violate the discrete maximum principle. We focus on a simplified version of the discrete maximum principle that provides only non-negativity of a discrete solution (the monotonicity condition.) This requirement is also very important since some physical quantities, such as concentration, temperature, or intensity, are nonnegative by their nature, and their approximations should be nonnegative as well. We propose the cell-centered finite volume method, which preserves nonnegativity of the continuous solution and demonstrates the second-order accuracy for the scalar unknown (Fig. 1).

The advection-diffusion equation can be written in the divergence form as

$$\operatorname{div} \mathbf{q} = f \text{ and } \mathbf{q} = \mathbf{v}c - \mathbb{K} \nabla c$$

where $\mathbb{K}(\mathbf{x})$ is a symmetric, positive, definite, continuous (possibly anisotropic) diffusion tensor, $\mathbf{v}(\mathbf{x})$ is a velocity field, f is a source term and $c(\mathbf{x})$ is an unknown concentration function. The vector function $\mathbf{q}(\mathbf{x})$ is an unknown concentration flux. The first equation represents the mass conservation law. The second one is the constitutive

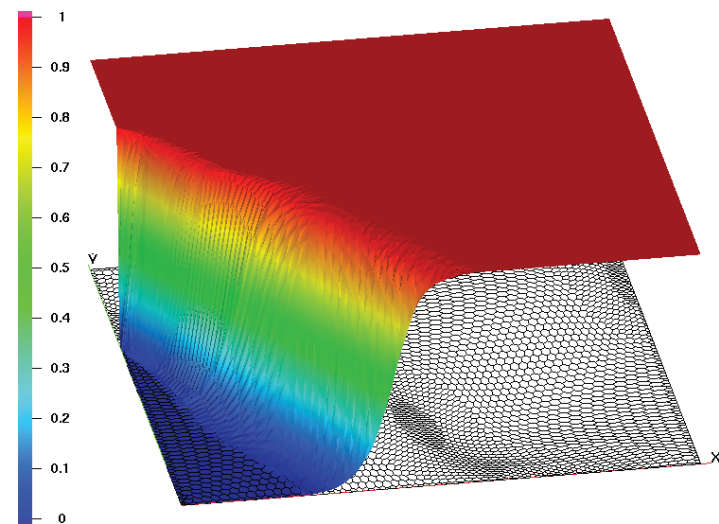


Fig. 1. Profile of the discrete solution on the polygonal grid. Problem: advection-diffusion equation with $Pe_K = 781,250$. The discontinuous Dirichlet-type boundary conditions result in the internal layer. The discrete solution satisfies the maximum principle and is between 0 and 1. The internal layer does not cause any numerical instabilities.

equation, which establishes connection between scalar and vector unknowns and is referred to as Darcy's law. An appropriate boundary condition should be imposed to make the problem well-posed.

The finite volume approach provides the local mass conservation by construction, since the equation is discretized in the mixed form. It means that two types of degrees of freedom are introduced: (1) cell-centered unknowns C_b for the unknown scalar function $c(\mathbf{x})$, and (2) unknowns \mathbf{q}_h for the normal component of the flux $\mathbf{q}(\mathbf{x})$ located on the edges. The mass conservation law can be easily discretized using Gauss-Green's theorem. The main difficulty, and the key to the discrete monotonicity condition, is a correct discretization of the constitutive equation. The concentration flux \mathbf{q} can be split into the diffusive flux \mathbf{q}_D , and the advective flux \mathbf{q}_A , which are defined as follows:

$$\mathbf{q}_D = -\mathbb{K} \nabla c \text{ and } \mathbf{q}_A = \mathbf{v}c.$$

The approximation of diffusive fluxes is based on a *nonlinear* two-point flux approximation method [1,2]. A *linear* two-point flux approximation has been known for a long time and results in the method that satisfies the discrete maximum principle. This feature makes this method very popular in many engineering applications. Unfortunately, in the case of anisotropic media and/or nonorthogonal grids, a discrete solution produced by this approach may be terribly inaccurate. The nonlinearity helps to rectify this problem by adopting the discretization method to a particular solution. The flux is also approximated at the middle of each edge using a weighted difference of concentrations in two neighboring cells. The nonlinearity comes from the fact that these weights depend on values of a discrete solution at cell centers in the neighborhood of this edge.

A proper discretization of the advective fluxes is especially important for advection-dominated problems. Those methods without special treatment of the advective fluxes produce nonphysical solutions with strong negative values and large overshoots. Here, we follow the ideas of the monotonic, upstream-centered scheme for conservation laws (MUSCL) introduced in [3]. In order to achieve second-order accuracy in the approximation, we use a piecewise linear discontinuous reconstruction of the finite volume (FV) solution on polygonal cells. To avoid unphysical oscillations in a numerical solution, a special slope-limiting technique is developed. Unlike in many existing approaches and similar to [3], we do not use a scalar limiter. Instead we limit a solution gradient by a matrix operator. Using this approach we reduce the limitation of a solution and obtain more accurate results. In each cell, we minimize deviation of the reconstructed linear function from given values at selected points subject to some monotonicity constraints. For each cell, the majority of these points are the centers of the closest neighboring cells, except in a few special cases.

The resulting nonlinear algebraic system is sparse, but nonsymmetric. For triangular meshes there are, at most, four nonzero elements in each row. We studied this system and proved that if there are no sinks in the system and the solution exists, then it is nonnegative. Moreover,

if the nonlinear system is solved using the Picard iterative method and the initial guess is nonnegative, then the proposed method guarantees nonnegativity of the discrete solution for all iterative steps. This means that for any tolerance of the iterative method we obtain a nonnegative solution. The proof is based on the special properties of M -matrices.

The numerical experiments presented in [4] demonstrate the flexibility and accuracy of the method in both advection-dominated and diffusion-dominated regimes. The method can be applied on unstructured polygonal meshes and full anisotropic heterogeneous diffusion tensors. The second-order convergence is observed for scalar unknowns C_b . For the normal components of the flux, q_b , the convergence is higher than the first order. The discrete solution remains nonnegative on different types of meshes and for different directions of anisotropy and velocity fields.

For more information contact Daniil Svyatskiy at dasyvat@lanl.gov.

-
- [1] C. Le Potier, *C.C.Acad. Sci. Paris, Ser. I* **341**, 787 (2005).
 - [2] K. Lipnikov, D. Svyatskiy, Y. Vassilevski, *J. Comput. Phys.* **228**, 703 (2009).
 - [3] M.E. Hubbard, *J. Comput. Phys.* **155**, 54 (1999).
 - [4] K. Lipnikov, D. Svyatskiy, Y. Vassilevski, LA-UR 09-03209. (submitted to *J. Comput. Phys.*)

Funding Acknowledgments

DOE Office of Science Advanced Scientific Computing Research (ASCR) program in Applied Mathematical Sciences.

Understanding the Nonlinear Physics of Laser-plasma Interaction through "At-scale" Plasma Kinetic Simulations on Roadrunner

Lin Yin, Brian J. Albright, Kevin J. Bowers¹ XCP-6; Ben Bergen, CCS-2

¹Guest scientist currently with D.E. Shaw Research, LLC, New York, NY 10036

In 2010, inertial confinement fusion (ICF) experiments commence at the National Ignition Facility (NIF). In these, over a million Joules of laser energy are focused within a gas-filled hohlraum. The hohlraum walls absorb the laser energy and re-radiate it as X-rays, which absorb in a spherical capsule at the hohlraum center. This causes the capsule to compress, bringing the deuterium-tritium fuel to the high temperatures and pressures required for thermonuclear fusion.

To prevent the hohlraum walls from ablating during the $\sim 10^{-8}$ s laser drive, a fill gas of hydrogen or helium is used. As the laser propagates through the fill gas, laser-plasma instabilities (LPI) may arise, which scatter laser light out of the hohlraum, degrade capsule implosion symmetry, and preheat the fuel with hot electrons, making compression harder to achieve.

Stimulated Raman scattering (SRS), the resonant amplification of electron density fluctuations by a laser, is one of the LPI concerns in ICF. In ICF experiments, a roughly uniform laser intensity is maintained across the beam with random phase plates that break the beam into an ensemble of laser speckles. For the success of fusion experiments on the NIF, we must first understand the physics of onset and saturation of SRS in the fundamental building block of a NIF laser beam, a single laser speckle. In a laser speckle, SRS manifests as the amplification of a forward-directed electron plasma wave (EPW) and the backward scattering of laser light. Unlike the linear growth of SRS, the nonlinear physics was not well understood

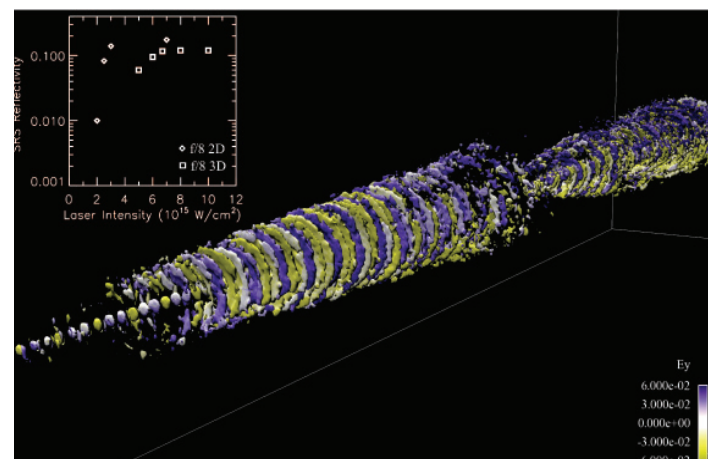


Fig. 1. Single-speckle LPI calculations using 16 CU of Roadrunner (11,520 ranks), nearly the full system; this calculation employed a record 0.4 trillion particles, >2 billion cells, and ran for nearly 58,160 time steps ($\sim 10^{19}$ floating point operations), long enough for two bursts of stimulated Raman scattering to grow from noise to significant amplitude at a laser intensity near the SRS onset. The inset shows the scaling of SRS reflectivity vs. laser intensity for 2D and 3D VPIC calculations; in 3D, the onset threshold is higher, owing to increased diffraction and side loss of hot electrons from laser speckles.

until recently [1,2]. Roadrunner has been used to assess the impact of the nonlinear SRS physics on laser penetration and energy deposition in fusion experiments. These fully kinetic plasma simulations employ the VPIC particle-in-cell code [3-6] and are performed in large plasma volumes in 3D at unprecedented range of time and space scales.

Until recently it has not been possible to simulate fully the comparatively large 3D plasma volumes of laser speckles. With VPIC on Roadrunner, simulations of the NIF hohlraum plasma have been done using 4096 IBM Cell Broadband Engine (Cell BE) chips at a range of laser intensity values (see the inset in Fig. 1). These simulations [7] showed that SRS reflectivity within a solitary speckle exhibits nonlinear behavior: a sharp onset at a threshold intensity, whereby reflectivity increases abruptly to a level orders of magnitude higher than linear theory predicts over a small range of intensity,

with a plateau in reflectivity at higher laser intensity in which SRS nonlinearly saturates; this generic behavior matches that measured in single speckle experiments at the LANL Trident Laser facility [8] with physics that cannot be captured by linear gain models of SRS growth within the speckle.

As a highlight of the unique simulations afforded by Roadrunner, the largest of these calculations was run on 16 connected units (CU) using 11,520 ranks, nearly the full Roadrunner system, and employed a record 0.4 trillion particles, over 2 billion computational cells, and ran for nearly 58,160 time steps ($\sim 10^{19}$ floating point operations), long enough for two bursts of stimulated Raman scattering to grow from noise to significant amplitude at a laser intensity near the SRS onset. Figure 1 shows isosurfaces of electrostatic field associated with these bursts; the wave fronts exhibit bending or bowing, arising from nonlinear electron trapping, as well as self-focusing, which breaks up the phase fronts.

The essential nonlinear physics governing SRS saturation has now been identified. The scattering manifests as a series of pulses, each of which passes through four distinct phases: 1) SRS grows linearly from density fluctuations, 2) electrons trapped by the EPW reduce the wave frequency and phase velocity by an amount that scales with EPW amplitude, 3) near the speckle center, where the amplitude is highest, the EPW phase velocity is lower than at the speckle's edge—EPW phase front bending ensues as shown by the top image in Fig. 2, and 4) the EPW wave amplitude exceeds the electron trapped particle modulation instability (TPMI) threshold [9,10]. TPMI generates waves off-axis from the laser direction and leads to EPW filamentation, self-focusing, and phase front breakup, shown by the bottom image in Fig. 2. Self-focusing increases transverse loss of trapped electrons and increases EPW damping.

From these basic science simulations, researchers are now able to better understand the essential nature of LPI nonlinear onset and saturation. Current research focuses on determining whether neighboring speckles can interact via exchange of hot electrons

or waves to produce higher backscatter than they would individually, the kind of study only possible on Roadrunner, where at-scale kinetic simulations of laser-plasma interaction in 3D at realistic laser speckle and multispeckle scales can be prosecuted at unprecedented size, speed, and fidelity.

**For more information contact
Lin Yin at lyin@lanl.gov.**

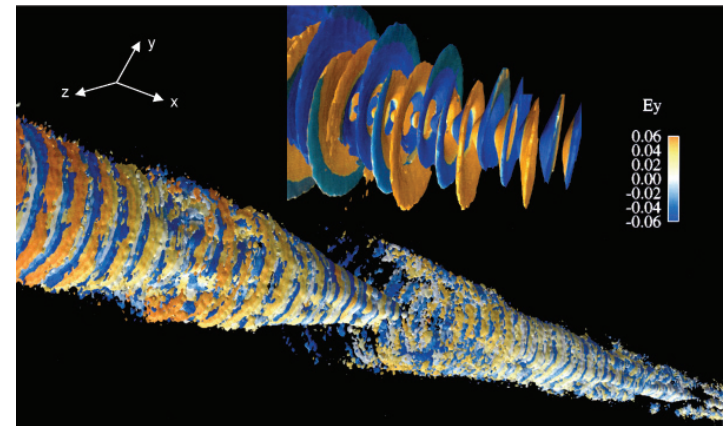


Fig. 2. (Top) A snapshot taken from a 3D VPIC LPI simulation at SRS pulse saturation of a $f/4$ laser beam, showing bending of iso-surfaces of EPW electric field across the speckle. The iso-surfaces are colored by the laser electric field intensity, graphically showing the source for SRS backscatter has become incoherent. **(Bottom)** Self-focusing and filamentation in two bursts of SRS, a snapshot taken from a 3D VPIC LPI simulation of a $f/8$ laser beam. The speckle volume is 16x larger than the $f/4$ simulation, permitting more transverse self-focusing modes to develop. This leads to chaotic EPW phase variation across the speckle. This further reduces SRS source coherence and increases wave damping, quenching the SRS pulse.

- [1] L. Yin et al., *Phys. Rev. Lett.* **99**, 265004 (2007).
- [2] L. Yin et al., *Phys. Plasmas* **15**, 013109 (2008).
- [3] K.J. Bowers et al., *Phys. Plasmas* **15**, 055703 (2008).
- [4] K.J. Bowers et al., *Proc. ACM/IEEE Conf. Supercomputing*, IEEE Press (2008).
- [5] K.J. Bowers et al., "Advances in Petascale Kinetic Plasma Simulation with VPIC and Roadrunner," *Proc. SciDAC* (2009).
- [6] K.J. Bowers et al., *J. Phys. Conf. Ser.* **180**, 012055 (2009).
- [7] L. Yin et al., *Phys. Plasmas* **16**, 113101 (2009).
- [8] D.S. Montgomery et al., *Phys. Plasmas* **9**, 2311 (2002).
- [9] H.A. Rose, *Phys. Plasmas* **12**, 012318-1 (2005).
- [10] H.A. Rose, L. Yin, *Phys. Plasmas* **15**, 042311 (2008).

Funding Acknowledgments

LANL Directed Research and Development Program

Robust Mesh Motion Based on Monge-Kantorovich Equidistribution

Luis Chacón, ORNL; Gian Luca Delzanno, John M. Finn, T-5

A method of grid adaptation using Monge-Kantorovich optimization has recently been developed [1]. This method enforces equidistribution of an error estimate of a discretized partial differential equation (PDE), based on the principle that the total error is minimized if the grid is chosen to equidistribute the local error. In 1D this condition determines the grid uniquely; in higher dimensions, equidistribution alone does not determine the grid uniquely. In Monge-Kantorovich theory, one minimizes

$$W \equiv \int_X \rho(\mathbf{x}) |\mathbf{x}' - \mathbf{x}|^2 d\mathbf{x}$$

enforcing equidistribution as a constraint. The variables $\mathbf{x}(\zeta)$ and $\mathbf{x}'(\zeta)$ are mapped from the logical or computational space Ξ , which is the unit square (unit cube in 3D) with a uniform grid, to the physical space X , which can be fairly arbitrary [2]. For time-stepping applications, the images $\mathbf{x}(\zeta)$ and $\mathbf{x}'(\zeta)$ for a uniform grid on Ξ represent the adapted grid on X at two different times. Equidistribution is enforced by requiring that the Jacobian $J \equiv \det(\partial x'_i / \partial x_j) = \rho(\mathbf{x}) / \rho'(\mathbf{x}')$, where ρ and ρ' are error estimates at the two times. Minimizing W with the equidistribution constraint leads to the result

$$\mathbf{x}' = \nabla \Phi(\mathbf{x})$$

and substituting this gradient condition into the equidistribution condition leads to the classic *Monge-Ampère equation* [1] or MA equation

$$\nabla^2 \phi + \phi_{xx} \phi_{yy} - \phi_{xy}^2 = \frac{\rho(\mathbf{x})}{\rho'(\mathbf{x}')} - 1$$

where $\Phi(\mathbf{x}) = \mathbf{x}^2/2 + \phi(\mathbf{x})$. Notice that nonlinearity enters in two ways: in the quadratic terms on the left (the Hessian), and in the

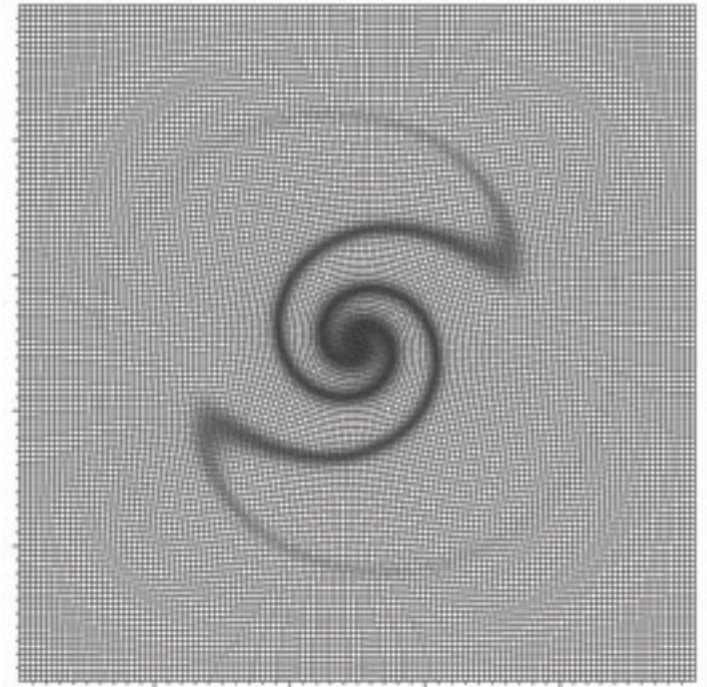


Fig. 1. Snapshot of the direct mesh at $t = 90$.

denominator on the right (where $\mathbf{x}' = \mathbf{x} + \nabla \phi$). Neumann boundary conditions $\hat{\mathbf{n}} \cdot \nabla \phi = 0$ on the boundary ∂X are used, to assure that boundary points move to boundary points. The operator on the left with these boundary conditions has a null space $\phi \rightarrow \phi + C$. Correspondingly, we must require that the right-hand side be in its range, i.e., $\int_V [\rho(\mathbf{x}) / \rho'(\mathbf{x}') - 1] dV = 0$. This is the nonlinear analog of the solvability condition $\int \rho dV = 0$ in solving the Poisson equation with Neumann boundary conditions. This nonlinear solvability condition, which depends on the map $\mathbf{x}'(\mathbf{x})$ at each Newton iteration, must be applied on each iteration for convergence. In [1] it was shown that, for monitor functions ρ and ρ' whose variation is moderate, solutions to the Monge-Ampère equation also minimize the mean distortion, defined as the integral over X of the trace of

the covariant metric tensor. Grid folding occurs when the distortion grows to a degree that the sides of a cell intersect. In this sense, grids obtained by the Monge-Kantorovich method are robust to tangling.

For time-stepping applications, this approach to grid adaptation can be accomplished in two distinct manners. The first is the *sequential* approach, in which ρ and ρ' are error estimates at two successive time steps. In the second method, the *direct* method, ρ is the error estimate at the initial time and ρ' is again the error estimate at the next time step. The reason these two approaches give different grids is that a composition of gradient maps is not a gradient map. (The inverse map, on the other hand, is a gradient map $\mathbf{x} = \nabla_{\mathbf{x}'}\Psi$, where Φ and Ψ are Legendre transforms.) We have tested both the sequential method and the direct method for several examples in order to determine the quality of the grids produced, specifically the robustness to distortion that leads to folding or tangling.

The PDE used was the passive scalar equation for a given incompressible flow

$$\frac{\partial \chi}{\partial t} + \mathbf{v} \cdot \nabla \chi = 0$$

Rather than solve a discretized version of this passive scalar equation for χ using an error estimate, we have solved it for several relatively simple flows by the method of characteristics and proceeded to equidistribute χ *itself*. Equidistributing χ does not test the linking of the moving mesh approach with an error estimate, which is a goal for future work. Rather, equidistribution of χ was chosen as an efficient way to test the relation between the grids determined by the sequential and direct methods, respectively. That is, for the sequential method, we have $\rho(\mathbf{x}) = \chi(\mathbf{x}, t)$ and $\rho'(\mathbf{x}) = \chi(\mathbf{x}, t + \Delta t)$; the direct method takes $\rho(\mathbf{x}) = \chi(\mathbf{x}, 0)$ and $\rho'(\mathbf{x}) = \chi(\mathbf{x}, t + \Delta t)$.

One of the example computations we have performed uses differential rotation $\Omega(r) = 16\Omega_0 \max[r(0.5 - r), 0]$ on $X = (0, 1) \times (0, 1)$, with $\Omega_0 = -0.1$. The grid obtained with the direct method is shown in Fig. 1 at $t = 90$. Figure 2 shows the equidistribution error, the mesh distortion, and the maxima of the

absolute values of the x- and y-components of the grid velocity. These quantities are shown for both the sequential method and for the direct method. Clearly, the direct method is better with respect to equidistribution and distortion, and also has somewhat smaller grid velocities. After $t \approx 90$, the sequential method begins to lose accuracy and soon breaks. The direct method continues well past this point, giving a smooth mesh even when the passive scalar structures are much smaller than can be resolved on the grid. Based on this example and several other examples, we conclude that the sequential method works fairly well for short times, but distortions in the grid eventually develop. The direct method appears to be completely robust, and in our view is the method of choice.

For more information contact John M. Finn at finn@lanl.gov.

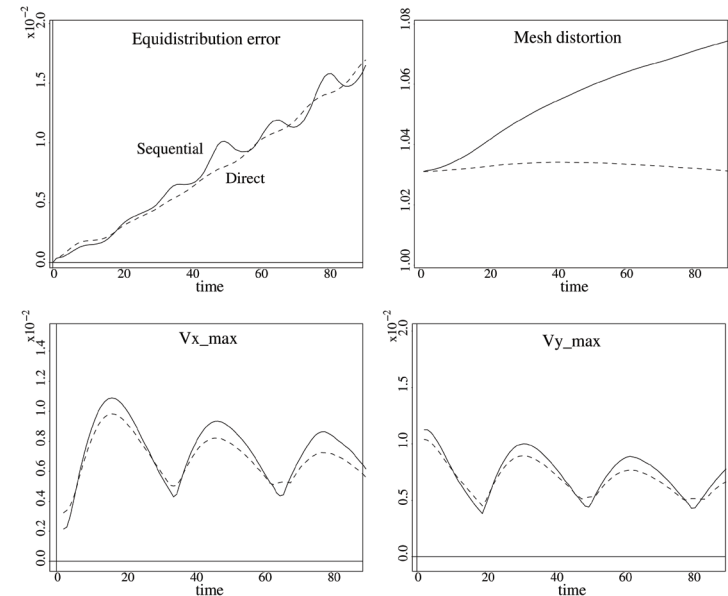


Fig. 2. Time histories of the equidistribution error, the mean mesh distortion, and the maxima of the absolute values of the grid velocity components.

- [1] G.L. Delzanno, L. Chacon, J.M. Finn, Y. Chung and G. Lapenta, *J. Comp. Phys.* **227**, 9841 (2008).
- [2] J.M. Finn, G.L. Delzanno, L. Chacón, *Proc. 17th Int. Meshing Roundtable*, Springer-Verlag (2008).

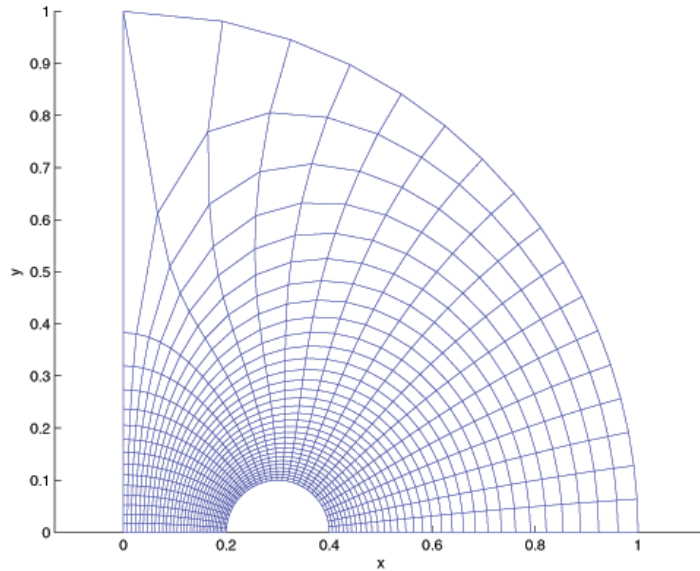
Funding Acknowledgments

LANL Directed Research and Development Program

A Particle-in-Cell Code with Arbitrary Curvilinear Mesh

Chris A. Fichtl, University of New Mexico; John M. Finn, T-5

Fig. 1. Physical space grid for studying the interaction of two dust grains in a collisionless plasma.



Particle-in-cell (PIC) codes solve the Vlasov equation, the basic equation of collisionless plasma kinetic theory. The Vlasov equation is a mean field equation, appropriate to hot, collisionless plasmas. That is, it is assumed that the particles interact only through the mean field, which is found by a global field solve from the particle positions and, for the case of an electromagnetic code, the particle velocities. The basic steps in an electrostatic PIC code are: 1) weighting particle positions to the grid, obtaining densities for the electron and ion species, 2) solving the Poisson equation for the electrostatic potential Φ on the grid and obtaining the electric field from Φ , and 3) integrating the particle orbits using this electric field, interpolated to the particle positions.

We are in the process of developing an electrostatic PIC code in 2D (cylindrical or axial symmetry) for an arbitrary structured mesh. The mesh is obtained by mapping from a uniform mesh on the *logical* (computational) domain Ξ , which we take to be the unit square, to an arbitrary non-uniform mesh on the physical space X . At present we are using the Winslow method [1] of mesh generation, a special case of the Laplace-Beltrami approach [2]. This method

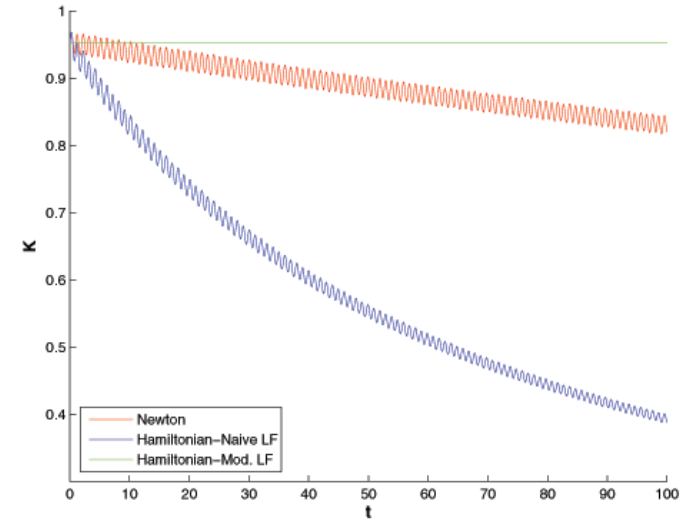


Fig. 2. The energy K as a function of time for three particle integrators.

produces boundary-conforming grids, allowing modeling of curved surfaces without the serious problems of "stair-stepping" in codes that use non-boundary-conforming grids. The example in Fig. 1 shows the mesh appropriate to dusty plasmas, in which the inner sphere corresponds to a dust grain, and the mirror image under a reflection through the vertical axis gives a second grain (the outer boundary represents matching to the rest of a Maxwellian plasma, and needs to be at a sufficiently large radius). Note that this method allows concentration of cells near a dust grain, to help in resolving the dynamics within a few Debye radii of the grain.

In our code all computations are performed on the logical grid on Ξ . When mapped to the logical grid the Poisson equation takes the form

$$\frac{\partial}{\partial \xi^i} \left(J g^{ij} \frac{\partial \Phi}{\partial \xi^j} \right) = -\frac{J \rho}{\epsilon_0}$$

(summation assumed) where $\rho \equiv \rho_{ion} - \rho_{electron}$ is the charge density, J is the determinant of the Jacobi matrix $J^i_j = \partial x^i / \partial \xi^j$, and $g^{ij} = \nabla \xi^i \cdot \nabla \xi^j$ is the contravariant metric tensor on the logical grid. The quantity $J\rho$ is the charge density on the logical grid.

The final hurdle is integrating the particle orbits (advancing or

pushing particles). The standard leapfrog integrator for a uniform grid is a symplectic integrator, i.e., conserves the Hamiltonian phase space structure. However, on a nonuniform grid standard (naive), leapfrog is not a symplectic integrator. We introduce a contact transformation $(\mathbf{x}, \mathbf{p}) \rightarrow (\xi, \mathbf{P})$ by a generating function $S(\mathbf{x}, \mathbf{P}) = \xi_j(\mathbf{x})P_j$, leading to a new Hamiltonian (for a time-independent grid)

$$K = \frac{1}{2m} \left(\frac{\partial \xi_j}{\partial x_i} P_j \frac{\partial \xi_k}{\partial x_i} P_k \right) + q\Phi(\mathbf{x}(\xi)) = \frac{1}{2m} \left(g^{jk}(\mathbf{x}(\xi)) P_j P_k \right) + q\Phi(\xi)$$

where q and m are the particle charge and mass, respectively. The fact that the Hamiltonian is not separable (it is of the form $K = T(P, \xi) + V(\xi)$) prevents naive leapfrog from being a symplectic integrator. To integrate the Hamiltonian equations of motion, we use a generalization (from 2D phase space to 4D) of the modified leapfrog integrator introduced in [3]. This time-centered, second-order accurate integrator involves a combination of implicit and explicit split-steps and preserves the Hamiltonian nature of the phase space exactly.

Figure 2 shows the orbits and a test of the energy conservation for this integrator on a nonuniform grid, compared with naive leapfrog integrators. The latter were performed with the straight Newtonian equations of motion and with the Hamiltonian equations using K , respectively. It is clear that in the nonsymplectic integrators, the orbits spiral in, losing energy, so that PIC runs must be severely limited in time step. The orbits from the symplectic integrator spiral without dissipation or growth for all time. Further, the energy is conserved to $O(\Delta t)^2$, the expected order in Δt and does not decay or grow systematically.

For integration into a full PIC code, an important factor is that, for uniform grids, the particle shapes required for density weighting and for electric field interpolation must be identical [4]. For nonuniform grids this condition must be modified: the two particle shapes must be *adjoints* of each other. If this condition is satisfied, the total momentum is conserved, and the so-called self-force is zero. Since our discretization is on the uniform logical mesh, the particle shapes are symmetric and this condition is automatically

satisfied, eliminating the need for computationally expensive particle search algorithms, which would be necessary if the computations were done on the physical grid.

We have performed the usual battery of tests to benchmark this code, including cold plasma oscillations, warm plasma Langmuir waves, Landau damping, and the linear and nonlinear two-stream instability, for a variety of nonuniform grids. For all, we have found good agreement.

**For more information contact
John M. Finn at finn@lanl.gov.**

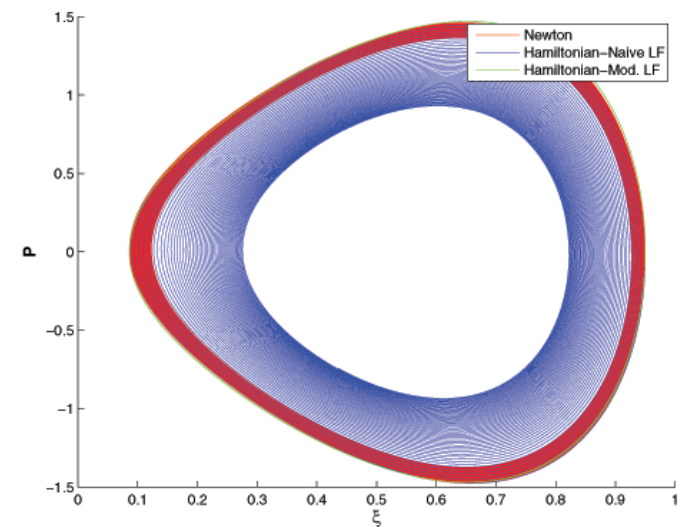


Fig. 3. Phase space for the three particle integrators, showing spiraling in for the two naive leapfrog integrators and good phase space conservation for the modified leapfrog integrator.

- [1] A.M. Winslow, "Adaptive mesh zoning by the equipotential method," *Tech. Rep. UCID-19062*, LLNL (1981).
- [2] V.D. Liseikin, *A Computational Differential Geometry Approach to Grid Generation*, Springer-Verlag (2007).
- [3] J.M. Finn, L. Chacón, *Phys. Plasma* **12**, 054503 (2005).
- [4] C.K. Birdsall, A.B. Langdon, *Plasma Physics via Computer Simulation*, Taylor and Francis Group (1995).

Funding Acknowledgments

DOE Office of Science

Robust Multigrid with Cell-based Coarsening by a Factor of Three

Joel E. Dendy, David J. Moulton, T-5

Multigrid methods are well established as one of the most efficient algorithms for the solution of the discrete linear systems that arise from models of diffusive phenomena (e.g., heat conduction, neutron diffusion, single-phase saturated flow). In particular, their solution cost grows only linearly with the size of the discrete system. These methods achieve this optimal scaling through the recursive use of successively coarser discrete problems (i.e., a sequence of coarse-grid discrete operators) in conjunction with smoothing on each level (e.g., a single Gauss-Seidel iteration on each level) to damp the highly oscillatory errors associated with each grid.

However, developing a robust multigrid algorithm with optimal algorithmic scaling for models with highly discontinuous and anisotropic coefficients is a significant challenge. Here a naive approach based on creating a hierarchy of discrete problems through rediscrization at coarser resolutions leads to a fragile algorithm with a convergence rate that depends on the magnitude of the jumps in the diffusion coefficient. Thus, a fundamental advance in multigrid methods was the development of *Black Box* methods that use the fine-scale discrete model to construct, through a variational principle, the successively coarser coarse-grid operators [1]. In fact, robust methods for both structured grids, such as Black Box Multigrid (BoxMG), and unstructured grids, such as Algebraic Multigrid (AMG), use this variational approach.

Despite this powerful technique, additional constraints on key elements of the multigrid algorithms may complicate its design. For example, it may be desirable to preserve the cell structure in the hierarchy of discrete operators. This desire often arises with cell-

based finite volume methods, and with cell-based adaptive mesh refinement. Unfortunately, a straightforward application of the BoxMG method [1], when coarsening by a factor of two, does not preserve the cell-based structure.

To preserve the cell-based structure and maintain a fixed operator complexity we proposed a variant of BoxMG that uses coarsening by a factor of three [2]. This cell-based coarsening by three is shown schematically in Fig. 1(top, left), where the coarse-grid cells are shown by the darker lines, and coarse-grid cell-center unknowns are shown as the darker circles. This coarsening strategy ensures that the cell-centered unknowns are nested on coarser grids and exposes interesting connections between BoxMG, smoothed-aggregation based AMG, and Multiscale Finite Elements. The key components of the new BoxMG algorithm follows.

- **Galerkin Coarse-Grid Operator.** Variational coarsening, employing the Galerkin coarse-grid operator, can be shown to be optimal in the sense that it minimizes the error in the range of interpolation. Moreover, this approach makes no assumption about the form of the coarse-scale model. This feature makes it ideal for problems with fine-scale spatial structure in the coefficients, as rediscrization fails to capture the influence of this structure at coarser scales.
- **Operator-Induced Interpolation.** We developed a generalization of the operator-induced interpolation methodology for coarsening by two, to coarsening by three. First we split the points into coarse and fine points, as shown in Fig. 1(top, right) with interpolation of coarse points given by the identity. The fine points are then split into two types: fine points that are embedded in coarse-grid lines (shown as small boxes in Fig. 1, bottom, left), and fine points that are in the interior of a coarse-grid cell (shown as small hexagons in Fig. 1, bottom, right). We then average the discrete operator at these fine-grid points to define the entries in the corresponding block matrices (4×4 and 2×2 systems) that are inverted to yield the interpolation weights.

• **Pattern Relaxation.** Although standard smoothing techniques, such as colored Gauss-Seidel are applicable here, in many situations these methods require an extra sweep of points at the top and right boundaries. To address this weakness we developed a new relaxation technique dubbed *pattern relaxation*. This smoother is motivated by the structure of the interpolation operator and uses a four-color block Gauss-Seidel scheme that corresponds to the blocks shown schematically in Fig. 1. An interesting aspect of pattern relaxation is its potential for coarsening by larger factors n , as is common in many multiscale methods.

To demonstrate the robustness of the proposed BoxMG method with respect to discontinuous diffusion and removal coefficients, different types of boundary conditions, and grid dimensions that are not optimal multiples of three, we considered a suite of test problems. For example, we considered a simple Poisson problem with favorable grid dimensions ($3m$ and $3m+2$), as well as unfavorable grid dimensions ($3m+1$), where m is an integer grid-size parameter. Within each of these categories the convergence rate was shown to be independent of m . To explore discontinuous diffusion coefficients we considered a thin layer problem as well as a variant of the classic checkerboard problem. Performance in these cases was similar to the simple Poisson problem with an average convergence factor for the largest problems of approximately 0.12. In addition, we demonstrated the effectiveness of the algorithm for mixed boundary conditions.

We have developed a robust multigrid method that uses coarsening by three to preserve the cell-based structure of the fine-grid problem. In this method we generalized the concept of operator-induced interpolation and developed a new relaxation method dubbed pattern relaxation. In the future we will investigate the extension of pattern relaxation to anisotropic problems. In addition, we will explore the connections of this new multigrid method with smoothed-aggregation based AMG and multiscale methods.

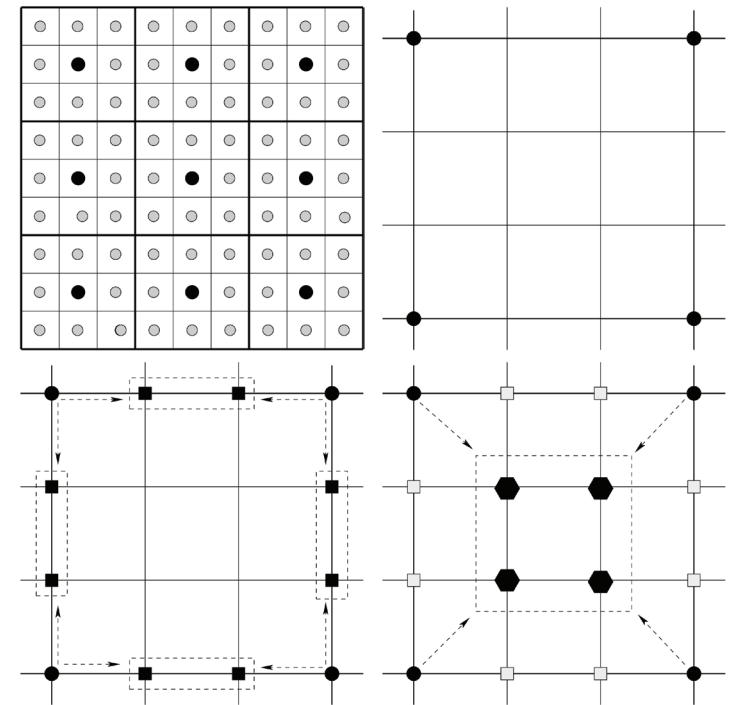


Fig 1. Cell-based coarsening by three is shown in the upper left. The structure of the operator-induced interpolation is shown schematically by the sequence of figures on a coarse-grid cell: injection of nested points in the upper right, followed by a block update of points embedded in coarse-grid lines (lower left) and finally the block update of fine-grid points contained in a coarse-grid cell (lower right). This operator-induced interpolation is critical to the method's robustness.

For more information contact David Moulton at moulton@lanl.gov.

- [1] J.E. Dendy, *J. Comput. Phys.*, **48**, 366, (1982).
- [2] J.E. Dendy, J.D. Moulton, *Numer. Linear Algebr.* (to appear) (2010).

Funding Acknowledgments

DOE, Advanced Scientific Computing Research (ASCR) program in Applied Mathematical Sciences

Numerical Simulation of Single Droplet Dynamics with Species Diffusion for Fuel Reprocessing

Marianne M. Francois, Neil N. Carlson, CCS-2

Fuel reprocessing employs liquid-liquid extraction processes to separate the components of used nuclear fuel. Two immiscible fluids (an aqueous phase and an organic phase) are counter-currently contacted, and the extraction process takes place at the liquid-liquid interface. The fluids are mixed together by mechanical action (e.g., rigorous stirring, rotation, pulsation), resulting in the dispersion of one phase into the other phase to ensure a large surface area between the phases, i.e., droplets of various size and size distribution are generated. Understanding the interface dynamics, droplet interaction, and extractant effect on droplet break-up and coalescence is very important as these small-scale phenomena can significantly affect the overall flow, and hence the extraction efficiency. Highly resolved interface dynamics simulations of the fluid flow, with species diffusion and reaction, will offer detailed information on the effects of fluid properties and flow regimes on mass transfer, drag coefficient, droplet size distribution, etc.—all required for the formulation of the reactive multiphase flow at the device scale.

To achieve this goal, we need to develop our simulation capabilities from our existing ones. As a first step, we have extended the TRUCHAS code [1] capability by coupling the current fluid-flow solver with interface tracking and surface tension [2] to the species diffusion solver [3]. The species advection term is added as a source term to the species diffusion equation. At the liquid-liquid interface, we assume thermodynamic equilibrium and assume the distribution coefficient to be unity for now.

To demonstrate the new capability, we have performed simulations in 2D and 3D of a single droplet rising by buoyancy with single species diffusion. The droplet is initially circular in 2D (spherical in 3D) in a rectangular domain, and initially contains the species. The gravitational acceleration is acting downward. We have investigated the effects of the Reynolds number on the droplet dynamics and species concentration (see Fig. 1). The Reynolds number represents the ratio between the inertial force and viscous force. As expected, we have found higher mass transfer with higher Reynolds number. We have also found our simulation results to be in agreement with the results of [4,5] for similar parameters. The droplet shape and species concentration contours for a 3D simulation are shown in Fig. 2.

Future work will include development of new capabilities to model chemical reaction, Marangoni effects, and surfactants.

For more information contact Marianne M. Francois at mmfran@lanl.gov.

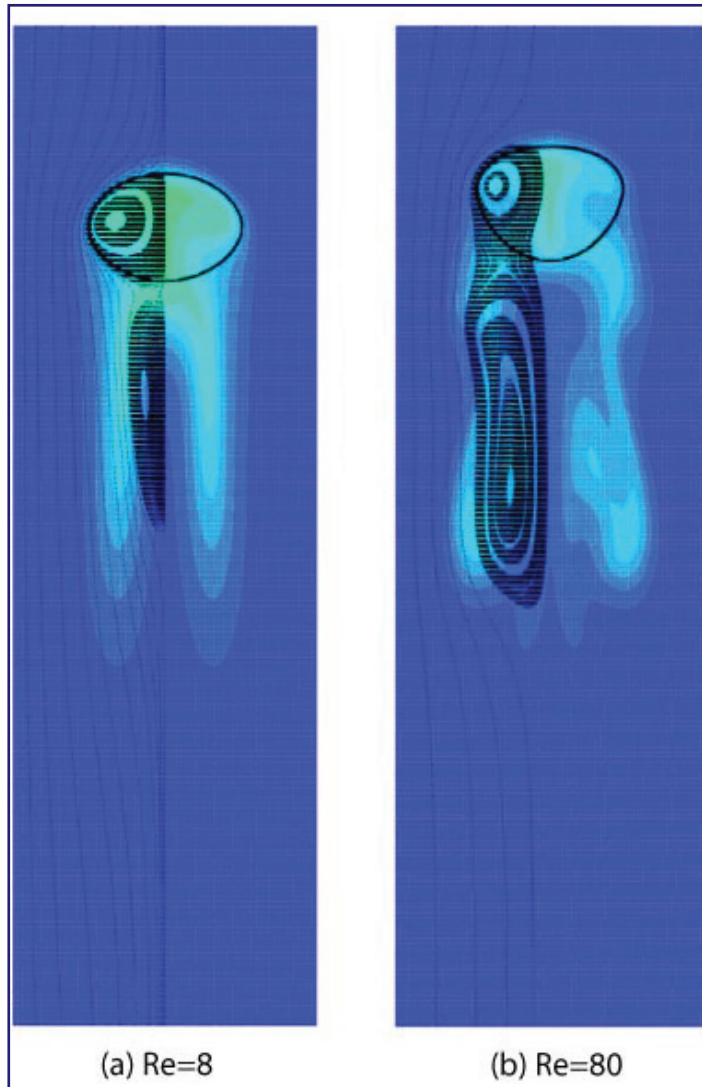


Fig. 1. Effects of Reynolds number (Re) on 2D droplet dynamics with Weber number of 0.22 and Péclet number of 1600. The species concentration contours and droplet shape with streamtraces (on left half) are plotted at time $t=1.8$ in a reference frame moving with the droplet. The Weber number represents the ratio between inertia and surface tension, and the Péclet number represents the ratio between advection and diffusion.

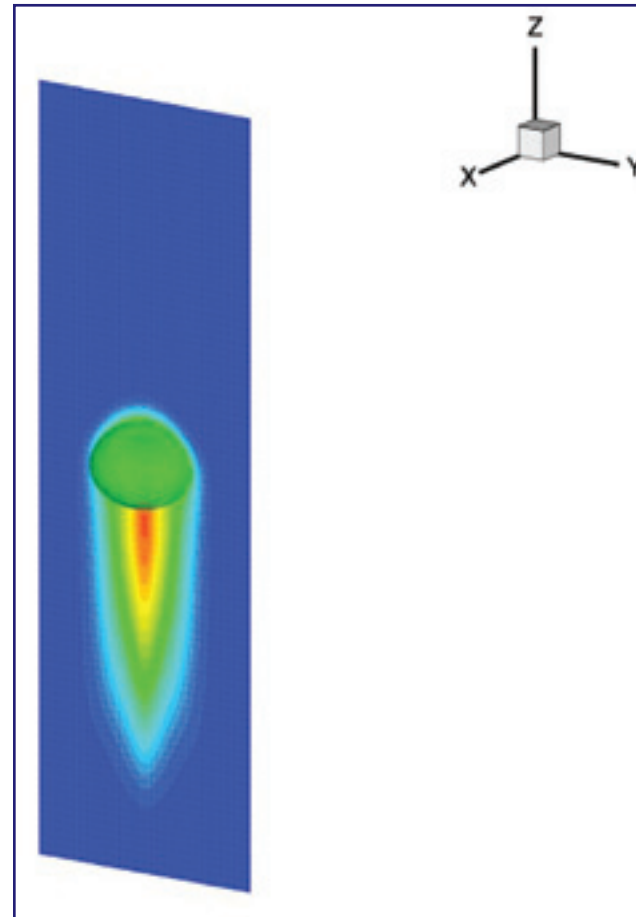


Fig. 2. Droplet shape and species concentration contours for 3D droplet simulation at $t=0.8$ in mid-plane. The Reynolds number is 8, the Weber number is 0.22, and the Péclet number is 1600.

- [1] <http://telluride.lanl.gov/>
- [2] M.M. Francois et al., *J. Comp. Phys.* **213**, 141 (2006).
- [3] N.N. Carlson et al., *PEDERNAL*, LA-UR 09-07050 (2009).
- [4] A. Koyonov et al., *AIChE J.* **51**, 2786 (2005).
- [5] Z.-S. Mao et al., *Int. J. Heat Mass Trans.* **44**, 1235 (2001).

Funding Acknowledgments

Nuclear Energy Advanced Modeling and Simulation program (NEAMS)

Astrophysics and Cosmology

A significant number of articles in this section exploit the tremendous computational resources offered by Roadrunner, the first petaflop/s machine in the world. Two articles are the result of work done in association with the Roadrunner Universe Project. The first describes the project itself; the second deals with baryon oscillations in the intergalactic medium. The work described in other articles used Roadrunner to study shocks in supernovae and National Ignition Facility experiments.

The evolution of the cosmos is the topic of an article that treats the universe as a black box. This is a problem where the researchers bring astrophysics, cosmology, and statistical science to bear to tackle one of the fundamental problems of cosmology. One article covers a study of models of the progenitors and the spectra of white dwarfs. The conductivity of neutron star crusts is the focus of research that uses molecular dynamics as the primary modeling tool.

The orbits of planets is the subject of an article on the effects of gas disks on the motion of the planets. The final article in the section reports on the origin of magnetic fields in galaxy clusters.

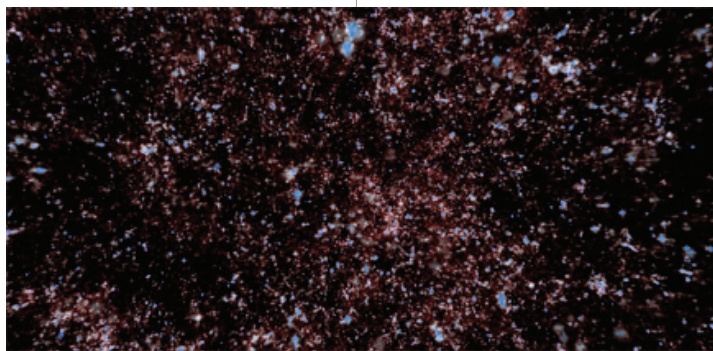
All these articles are representative of the broad interests of ADTSC staff in the fields of astrophysics and cosmology.

Astrophysics and Cosmology

Hybrid Petacomputing Meets Cosmology: The Roadrunner Universe Project

Salman Habib, T-2; Adrian Pope, ISR-1/T-2/CCS-6; Zarija Lukić, T-2; David Daniel, CCS-1; Patricia Fasel, CCS-3; Katrin Heitmann, ISR-1; Nehal Desai, CCS-1 and Aerospace; Chung-Hsing Hsu, CCS-1 and ORNL; Lee Ankeny, HPC-1; Graham Mark, CCS-3; Suman Bhattacharya, T-2; James Ahrens, CCS-1

Fig. 1. Dark matter halo visualization from a fly-by zoom-in view of a large cosmological simulation with MC³ (visualization by J. Woodring, CCS-1). Galaxies inherit their clustering properties from those of the host dark matter halos in which they reside.



Over the last two decades, critical observational advances in large-volume sky surveys carried out over a wide range of wavelengths, as well as over short time cadences, have revolutionized cosmology. Computational cosmology has emerged as an essential resource for providing detailed predictions for these observations, essential data for assisting in their design, and sophisticated tools for interpreting the final results.

Results from cosmological surveys have cemented a cross-validated cosmological Standard Model, presenting a comprehensive picture of the evolutionary history of the Universe and its constituents: 23% in dark matter, which only interacts gravitationally (and a large fraction of which is in localized clumps called halos), and 72% in a smooth dark energy component that is described by a cosmological constant, adiabatic Gaussian random initial density fluctuations, and flat spatial geometry [1]. Although this result is a great triumph, it has exposed some of the biggest puzzles in physical science: What is dark

matter? Why is the expansion of the Universe accelerating? Does general relativity need to be modified? What is the origin of primordial fluctuations?

To investigate these questions, the observational state of the art is rapidly advancing; surveys now coming online and within the next decade represent an

improvement in capability by roughly two orders of magnitude, translating into a determination of certain cosmological parameters at the 1% level. Remarkable as this is, the effort will only come to fruition if the accuracy of the underlying theory can be controlled to the sub-percent level. This severely demanding task will push the boundaries of computing for the foreseeable future.

With one or two exceptions, all cosmological probes are based on statistical studies of fluctuations, whether temperature in the case of the cosmic microwave background (CMB), or of the mass distribution in the case of weak gravitational lensing. The fundamental task for theory is to produce accurate predictions for the CMB, for the mass density and velocity field (along with associated statistical measures such as the fluctuation power spectrum), and for tracers of mass and velocity, such as galaxies and galaxy clusters. Cosmological information on large spatial scales is relatively easier to interpret because on these scales the physics is essentially linear—however, observable quantities can have a large variance due to finite sampling limitations. At smaller, more nonlinear scales, the statistical limitations can be essentially removed, but modeling becomes significantly more complicated, so keeping the associated systematic errors in check is a difficult challenge. There is currently no alternative to precision simulation as the theoretical tool of choice for dealing with the nonlinear regime of structure formation.

Structure formation in the Universe is driven primarily by the gravitational instability. Initial density perturbations collapse and merge in a hierarchical fashion to form dark matter halos within a global *cosmic web* structure (Fig. 1). On scales smaller than several megaparsecs (Mpc) (1 parsec = 3.26 light-years), baryonic matter collects in halos, eventually forming stars and galaxies. The collisionless evolution of matter subject only to gravity is described by the Vlasov-Poisson equation in an expanding Universe, which can be solved in detail only by N-body techniques. Next-generation surveys demand simulations with multi-gigaparsec (Gpc) box-sizes and particle counts in the 10^{11-12} range, all with approximately kiloparsec (kpc) force resolution (a force dynamic range of 10^6). An overall two to three orders of magnitude improvement in

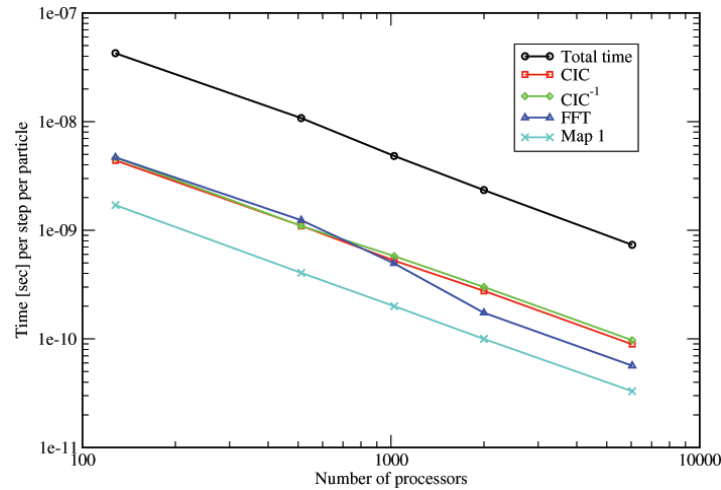


Fig. 2. Scaling of wall-clock time per particle per step as a function of the number of processors. The total problem size is scaled roughly with the number of processors (weak scaling).

throughput over the current state of the art turns out to be the minimal requirement.

To meet the challenge of next-generation simulations, the Roadrunner Universe (RRU) project at LANL, using the IBM Cell Broadband Engine (Cell BE), has developed the Mesh-based Cosmology Code on the Cell BE (MC³), the first hybrid petascale cosmology code. The MC³ algorithm splits the inter-particle force problem into two parts, a medium-resolution solver based on fast Fourier transforms (FFTs), augmented by a direct particle-particle short-range solver. The biggest FFT provides up to four orders of magnitude of dynamic range, the remaining factor of 10–100 coming from the short-range force evaluations carried out on the Cell BE processors of Roadrunner. The MC³ algorithms match to the machine architecture, minimizing data transfer through the narrow communication pipe between the central processing units (CPU) and the Cell BE. The global philosophy is to explicitly sacrifice memory and in-place computation in order to minimize communication and simplify communication patterns.

Our approach has two key aspects: 1) reduction of particle communication across the Cell layer using particle overloading, a 26-neighboring processor-mirrored particle cache, and 2) application of digital filtering and differencing in the spectral domain, allowing simplified computations at the Cell BE layer [2]. While MC³ was originally developed for Roadrunner, the computational strategy can easily be modified for other petascale platforms because the short-range and long-range force solvers are effectively decoupled, allowing for independent optimizations. As a result of its design, MC³ possesses excellent parallel scaling properties (Fig. 2). Since the size of the FFT is quite manageable for current and next-generation supercomputers, MC³ will scale to the largest problems that can be run on these machines. Additionally, the MC³ implementation is likely to evolve smoothly as future architectures add further layers of memory and computational hierarchies.

The first scientific application of MC³ is an analysis of the baryon acoustic oscillation (BAO) signal in the quasar Lyman- α forest [3]. Fixing the BAO scale via cosmological measurements is probably the least systematic-affected technique for investigating the dark energy equation of state. The work involved running nine of some of the largest cosmological simulations ever performed (64 billion particles) with a dynamic range sufficient to resolve the smallest scale of interest (the Jeans scale), as well as enough volume to realistically capture the depth of the Baryon Oscillation Spectroscopic Survey (BOSS), the key cosmological component of the Sloan Digital Sky Survey III. BOSS aims at a percent-level measurement of the BAO scale.

An extensive cosmological simulation program is planned for MC³, including applications to next-generation surveys such as the Large Synoptic Survey Telescope (LSST) project along with a variety of physics extensions, such as self-consistent dark energy models, modified gravity, non-Gaussian initial perturbations, and gas physics.

[1] E. Komatsu et al., *Astrophys. J. Supp.* **180**, 330 (2009).

[2] A. Pope et al., *Comp. Sci. Eng.* (to appear).

[3] M. White et al., *Astrophys. J.* arXiv:0911.5341 (to appear).

Funding Acknowledgments

LANL Directed Research and Development Program

For more information contact
Salman Habib at
habib@lanl.gov.

Cosmic Emulation: The Universe as a *Black Box*

Earl Lawrence, CCS-6; Katrin Heitmann, ISR-1; David Higdon, CCS-6;
Christian Wagner, ISR-1; Martin White, University of California, Berkeley;
Salman Habib, T-2; Brian Williams, CCS-6

Cosmology is perhaps the grandest of inverse problems: given knowledge of observational data, we wish to infer the fundamental laws governing the dynamics of the Universe and the properties of its constituents. The difficulty, of course, is that the Universe is very complicated and very large, so that the cosmic inverse problem is by no means easy to solve.

However there is one regime where one can do well, and that is studying the Universe on very large scales. For example, the primary temperature anisotropies in the cosmic microwave background (CMB) are small and can be measured accurately on angular scales large enough so that a linear treatment of cosmological fluctuations suffices to describe them. As a result, the associated inverse problem – determination of a handful of cosmological parameters – can be solved in a relatively straightforward manner.

The story becomes much more complicated when tracers of the distribution of matter are involved and when one goes to small length scales where the physics becomes nonlinear and complex. There are two reasons why this ostensibly more complex regime is nevertheless of very significant interest. First, investigations of the structure of the Universe on large scales suffer from the fact that the associated statistical error bars are large due to unavoidable finite sampling limitations, the so-called *cosmic variance* problem. Fluctuations on small scales are much better sampled, with negligible cosmic variance. Second, it is very important to have multiple probes that can be crosschecked with each other in order to avoid systematic biases, and to include information missing from a certain cosmological probe, or class of probes, due to parametric degeneracies.

Solving the cosmic inverse problem has gained significant urgency with the discovery that the expansion of the Universe is actually accelerating, rather than slowing down. Either a very strange new component is driving this acceleration, or general relativity – our cherished theory of gravitation – may need to be modified. It is widely thought that cosmological measurements (and their interpretations) at the percent level of accuracy are required to begin to answer some of these questions. Because of the severe computational difficulties, as well as scientific uncertainties, large-scale numerical simulations and phenomenological approaches are the only tools available to address the inverse problem.

The classic approach to a statistical inverse problem is via Bayesian inference, where the posterior distribution is sampled by Markov Chain Monte Carlo (MCMC) methods. This approach involves solving the forward model tens to hundreds of thousands of times. If the forward model is a complicated numerical code that takes days to run, the MCMC approach becomes impractical. To overcome this problem, we have introduced a framework that we call *cosmic calibration* [1,2]. An essential aspect of our approach is the use of interpolation in high dimensions, via principal components and Gaussian process modeling, to generate forward model predictions while having run only a small number of base simulations (of the order of tens to hundreds). We refer to the interpolation system as a forward model *emulator*.

A high-accuracy demonstration of this approach is presented in a sequence of three papers targeting precision emulation of the mass fluctuation power spectrum down to scales small enough to be relevant to current weak gravitational lensing surveys mapping the distribution of matter. The first paper [3] shows that the underlying simulations are accurate to the 1% level, the second [4] shows that the emulation scheme is also as accurate (by using an approximate stand-in model for the actual simulations), and the third paper [5] presents the full emulator. The remarkable aspect of having an operational emulator is that it replaces simulations of the Universe by an effectively infinitely fast *black box* oracle. Using the oracle can reduce the MCMC analysis times from years to minutes.

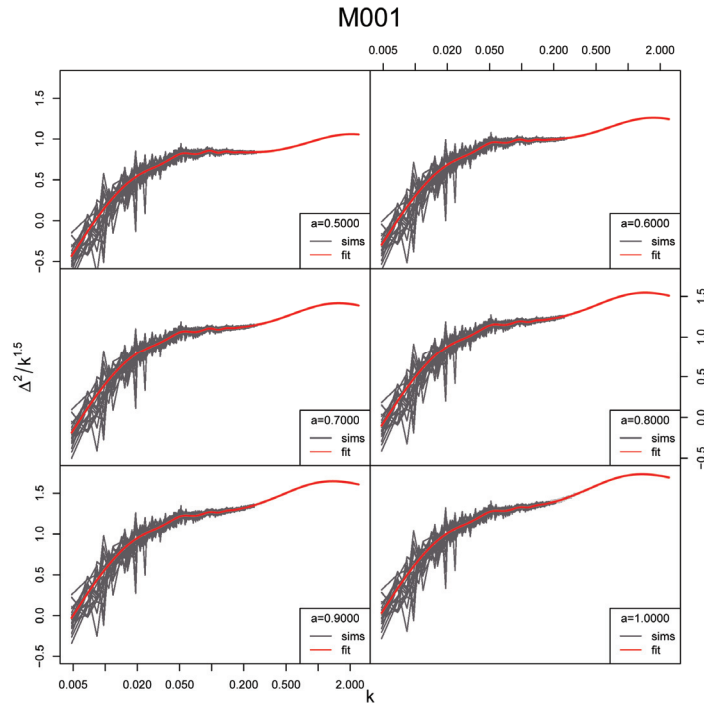


Fig. 1. Simulated dimensionless power spectra for the mass fluctuations in one of the 37 simulated cosmological models. The large fluctuations at small values of the wave number are due to cosmic variance, perturbation theory is used to fit the $k < 0.03 \text{ Mpc}^{-1}$ regime. At larger values of k , process convolution is used to smooth the simulation data—note that subtle features such as the small baryon acoustic oscillation wiggles are nicely fit by this technique.

The construction of the cosmic *black box* was accomplished by running a suite of multiresolution simulations on LANL's Coyote cluster. Almost 40 cosmologies were run with roughly 20 simulations per cosmology, for a total of almost 1000 large-scale simulations and an associated dataset size of 60 terabytes (TB), the largest simulation suite of its kind. The parameters for the cosmological models were sampled using Latin Hypercube sampling (for details, see [4]) over a prior range consistent with current observational constraints. Figure 1 shows the dimensionless power spectrum for one of the simulated cosmologies at six values

of the expansion factor. The smooth fit is obtained by combining perturbation theory at large scales and adaptively filtering the noisy simulation results at smaller scales using process convolution. Figure 2 demonstrates the accuracy of the emulation technique. Although only 37 sampling points were used over a five-dimensional space, the emulation accuracy for predicting a function over more than two decades of dynamic range is (mostly) better than a percent. We have made the emulator publicly available so that it can be used by the entire community.

Future efforts will be targeted to increasing emulation dynamic range, adding more cosmological parameters (dimensions), and constructing emulators for other observables such as the weak lensing shear power spectrum and the cluster mass function. Although our work has shown that emulation is a powerful technique, building successful emulators is technically challenging and requires a concerted effort among different communities.

For more information contact Salman Habib at habib@lanl.gov.

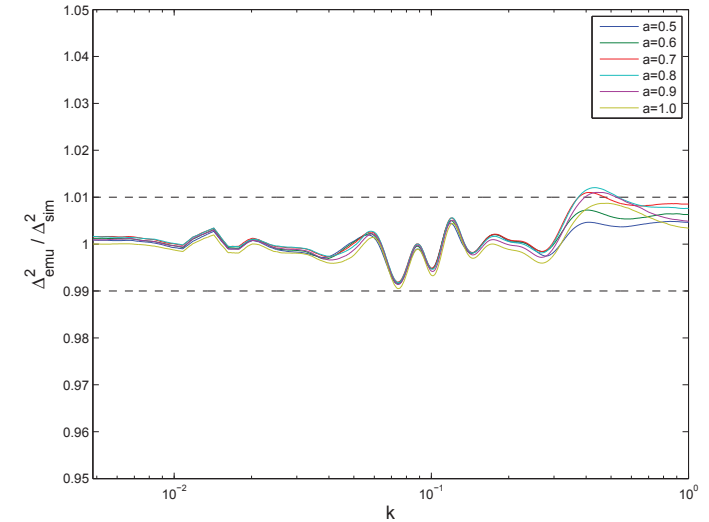


Fig. 2. Ratio of the emulator prediction to the smoothed simulated power spectrum for a Λ CDM cosmology at six values of the scale factor. The error exceeds 1% very slightly in only one part of the domain for the scale factor values, $a = 0.7, 0.8$, and 0.9 .

- [1] K. Heitmann et al., *Astrophys. J. Lett.* **646**, L1 (2006).
- [2] S. Habib et al., *Phys. Rev. D* **76**, 083503 (2007).
- [3] K. Heitmann et al., *Astrophys. J.*, arXiv:0812.1052 (to appear).
- [4] K. Heitmann et al., *Astrophys. J.* **705**, 156 (2009).
- [5] E. Lawrence et al., *Astrophys. J.*, arXiv:0912.4490 (to appear).

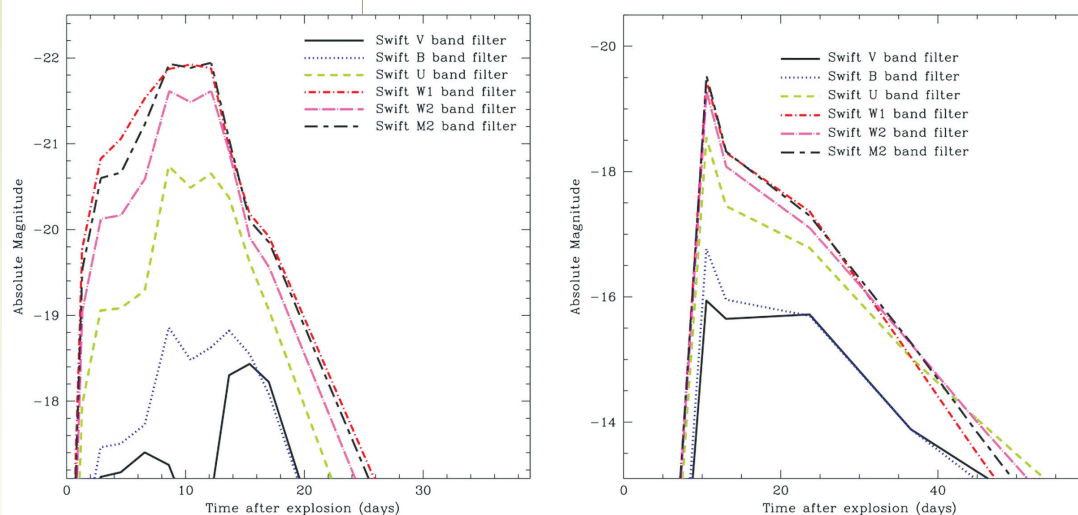
Funding Acknowledgments

- LANL Directed Research and Development Program
- NASA Theoretical Astrophysics Program
- DOE High Energy Physics

Modeling Progenitors and Spectra from White Dwarf Collisions: Thermonuclear and Core-Collapse Supernovae

Chris L. Fryer, CCS-2; Aimee Hungerford, XTD-6; Gabriel Rockefeller, CCS-2;
Bob Greene, XCP-2; Frank Timmes, Cody Raskin, Arizona State University

Fig. 1. Light-curves (luminosity as a function of time in astronomy magnitude units) for a range of Swift band filters [1]. A key feature of accretion-induced collapse explosions is that the ultraviolet emission is much stronger than the optical emission. The left plot shows a strong explosion exemplifying the strongest explosions produced for these objects. The right plot shows a weaker explosion. Both are plausible given the current set of error bars.



The past decade has seen a renaissance in observation strategies designed to detect astrophysical transients. The excitement garnered from the rapid gain in understanding of gamma-ray bursts spilled over to the transient field. Now over 20 fast-slewing telescopes exist in the worldwide astrophysics community to study transients. More importantly, large surveys with cadences designed to detect astrophysical transients have shown that the canonical picture outlining just a few classes of cosmic explosions oversimplified the diverse menagerie of objects now known as astrophysical transients. These first tantalizing observations have revived the study of failed or lost cosmic explosions: outbursts that didn't fit within the canonical supernovae.

One such class of lost explosions has been the outbursts produced in the collisions or mergers of two white dwarfs. Here we discuss two types of explosions produced by such collisions: 1) the gravitational-wave-induced merger of a stellar binary consisting of two white dwarfs whose total mass exceeds the Chandrasekhar limit, and 2) the collision of two white dwarfs. If two white dwarfs merge and the combined mass of this system is above the Chandrasekhar limit, electron degeneracy pressure can no longer support the merged object and it collapses. Scientists had predicted that such a system would collapse to form a neutron star and a weak supernova explosion. But until this past year, none had been observed, or if they had been observed, they were thrown out of any published sample. However, due to broader-interest surveys in astrophysical transients, observations were reported this year of a couple potential accretion-induced collapses, as these merger-induced collapses are termed.

Figure 1 shows the first detailed light-curve calculations of the collapse of a merged double white dwarf binary system [1]. We discovered that the light curves depend very sensitively on both the surrounding medium and the exact nature of the explosion. In working with observational astronomers in the Swift consortium, we produced light curves for each of the Swift bands. The ultraviolet provides an ideal diagnostic of the supernova type, and Swift observations will be critical in identifying the class of each observed outburst. We have also worked closely with the Large Synoptic Survey Transient team and calculated light curves for likely filters for this telescope (presented at the January 2010 American Astronomical Society (AAS) meeting). With these models, we were able to rule out some of the potential accretion-induced collapse candidates, but the identification of the remaining observations as white dwarf explosions is now on much more solid ground. Accretion-induced collapse outbursts may have very strong gravitational wave signatures (strongest of any stellar collapse event) and exhibit pristine neutrino signals. As such, they are ideal transients for studying general relativity, nuclear, and particle physics. Our preliminary work places us in an ideal position to dominate this field.

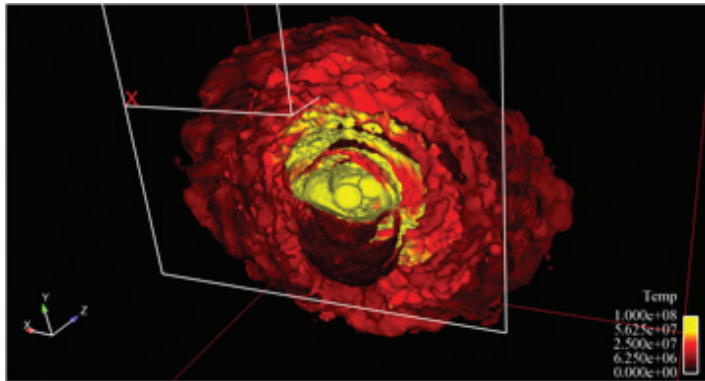


Fig. 2. Peering into the heart of two colliding white dwarfs. The color-coding denotes temperature. Shock heating can ignite a thermonuclear explosion, causing a supernova-like outburst.

With transient surveys discovering new classes of outbursts, scientists have also begun to study new ways to produce transient-like phenomena. Using LANL's SNSPH, a parallel 3D smoothed particle radiation hydration code, scientists at Arizona State University modeled the collision of two white dwarfs [2]. Figure 2 shows an image of this collision at a single point in time. In some cases, the collision is violent enough to ignite the white dwarf, driving a thermonuclear explosion. Although very different than the canonical type Ia supernovae (also known as thermonuclear supernovae), it demonstrates that thermonuclear explosions of white dwarfs can occur in a variety of manners. It may also explain a set of abnormal type Ia supernovae. Again, observations have discovered that nearly 30% of thermonuclear supernovae are not normal, and it is very possible that some of these abnormal supernovae arise from collisions.

Figure 3 shows the spectra produced by LANL's supernova light-curve effort. The LANL effort takes advantage of both codes and physics capabilities developed under the Advanced Simulation and Computing Program (ASC) and the Campaign programs at LANL, allowing this work to produce the first ever light curve (luminosity versus time) and corresponding detailed spectra (luminosity versus

photon energy) using a single two-temperature radiation hydrodynamics calculation. The detailed atomic opacities used in these calculations allowed us to produce detailed spectral images. The spectra in Fig. 3 can be compared with observations of these outbursts and distinguish collision-induced thermonuclear supernovae from the more canonical type Ia supernovae. Such work is critical for the Joint Dark Energy Mission's success. The Joint Dark Energy Mission assumes that we can use type Ia supernovae as standard candles (we know the intrinsic luminosity of an object so can determine its distance based on its apparent luminosity). If objects such as these colliding white dwarfs are infiltrating our high-red-shift sample of supernovae, they may skew the resultant calculation for dark energy parameters. We have been tasked as theorists to determine diagnostics so that we can exclude these anomalous outbursts from our sample.

These failed or lost white dwarf explosions have implications for a wide range of physics and astrophysics projects from the Joint (NASA/DOE) Dark Energy Mission to the NSF Laser Interferometer Gravitational-Wave Observatory led by general relativists to nuclear physicists working on the Facility of Rare Isotope Beams and the Deep Underground Science and Engineering Lab. These initial studies, leveraging LANL's computational and physics expertise developed under ASC and the Campaigns, place LANL at the forefront of these studies, and we are in an ideal position to take our initial successes and dominate this broad-implication research.

**For more information contact
Chris Fryer at fryer@lanl.gov.**

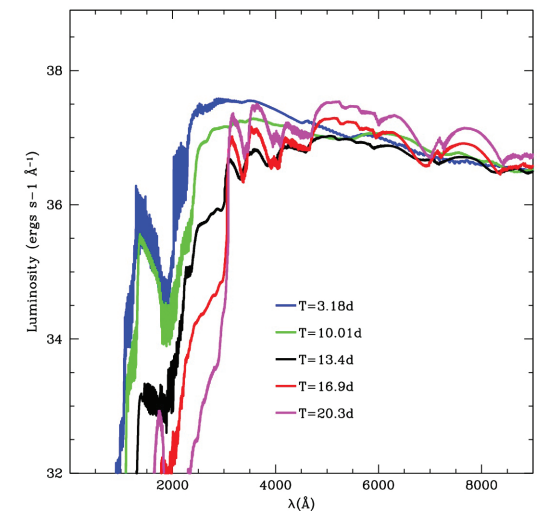


Fig. 3. Spectra from the outburst produced in the collision of two white dwarfs. These spectra were produced using the entire 14,900 group opacity data produce at LANL using the TOPS code and all the features are due to lines in the atomic physics. The more robust features can be used to distinguish these explosions from other explosions.

[1] C.L. Fryer et al., *Ap. J.* **707**, 193 (2009).

[2] C. Raskin et al., *Mon. Not. R. Astron. Soc.* **399**, L156 (2009).

Funding Acknowledgments

- LANL Directed Research and Development Program
- DOE Office of Science
- National Science Foundation (NSF) Theoretical Research

Molecular Dynamics Simulation of Multispecies Neutron Star Crust Conductivity

Sanjib Gupta , Jerome Daligault, T-5

Accurate knowledge of the thermal properties of the accreting neutron star crust, required to model the time evolution of temperature at all depths in the neutron star crust, is crucial to understanding their nonexplosive (quiescent) X-ray emission characteristics, and thereby to unravel the nature of exotic matter in the crust and core of these stars. Apart from the location and strength of the heat sources from various nuclear reactions [1,2], the most important determinant of crust thermal profile is the electron thermal conductivity of the crust. The thermal conductivity of the crust is governed by the electron scattering processes of the ionic centers, impurities, and defects, together with electron interactions with lattice vibration modes.

Calculation of the conductivity at different depths in a neutron star has been complicated by 1) the changing composition due to nuclear reactions, and 2) the uncertainty in the state (amorphous vs. crystalline) of a multispecies, strongly coupled mixture of nuclei with from a few (~ 10) to more than 300 species of varying nuclear charge Z and mass A ($5 < Z < 55$ and $10 < A < 135$ roughly). Whereas several analytical approaches have been applied to binary mixtures, their extrapolation to 300+ species has never been investigated or verified. A typical example is the so-called impurity parameter formalism in which a perfect crystal with a low concentration of impurities substituted at lattice sites can be approximated as being made of two uncorrelated parts—the perfect lattice together with the residual charge. The total electron-ion-scattering cross section is the sum of the Bloch electron-lattice cross section plus the sum of uncorrelated binary electron-residual

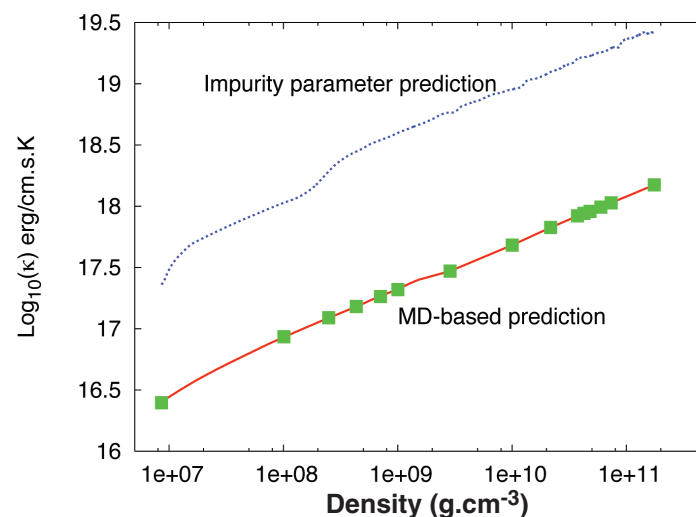


Fig. 1. Thermal conductivity as a function of density (g/cc) in the crust of an accreting neutron star. The MD calculations (green squares) are accurately reproduced by the linear mixing rule (red solid line) and are contrasted with the electron-impurity scattering conductivity (blue dotted line) over a range of four orders of magnitude in crust density. The MD conductivity is also less sensitive than the impurity parameter conductivity to sudden crustal composition changes from electron capture reactions triggered by density changes.

impurity collisional cross sections. However, as the concentration of impurities increases, this splitting is less and less justifiable. Indeed, by performing large-scale parallel molecular dynamics (MD) simulations of the charge density fluctuations (structure factor) of complex mixtures of nuclei, and over a range of densities spanning four orders of magnitude, we have shown in [3] that the impurity parameter formalism yields crust conductivities that are an order of magnitude higher than they should be (Fig. 1). This result has enormous implications for upcoming experimental campaigns that target X-ray transients.

Furthermore, we have verified a microscopic version of the linear mixing rule for the calculation of the structure factor of a complex mixture of nuclei that uses the structure factor of individual species

rather than quantities such as conductivity, which are integrated over the full range of momentum transfer in electron-ion collisions. This new mixing rule, which can be applied to mixtures of arbitrary complexity (see Fig. 2 for an example), can be used not only for calculating the thermal conductivity but also for equation of state and neutrino loss rates, thus drastically reducing the complexity of large astrophysical codes that simulate the thermal structure of neutron stars. We foresee much application of our new mixing rule to the prediction of X-ray emissions from transiently accreting neutron stars that will be observed in the future.

For more information contact Sanjib Gupta at guptasanib@lanl.gov.

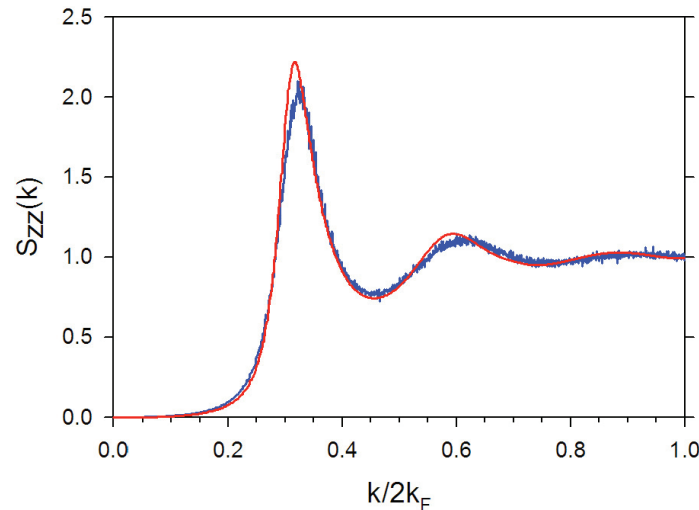


Fig. 2. The structure factor obtained with MD simulations (red line) and with the mixing rule formula (blue line) for a realistic mixture of 451 nuclear species at density $2.49 \cdot 10^8$ g/cc and temperature 0.5 billion Kelvin in the accreting neutron star crust. This newly discovered microscopic linear mixing rule, when expressed as a simple abundance-weighted average of individual species charge-charge structure factors, can be applied to a wide range of thermal quantities such as Equation of State (EOS) and neutrino loss properties.

- [1] S. Gupta et al., *Astrophys. J.* **662**, 1188 (2007).
- [2] S. Gupta et al., *Phys. Rev. Lett.* **101**, 231101 (2008).
- [3] J. Daligault, S. Gupta, *Astrophys. J.* **703**, 994 (2009).

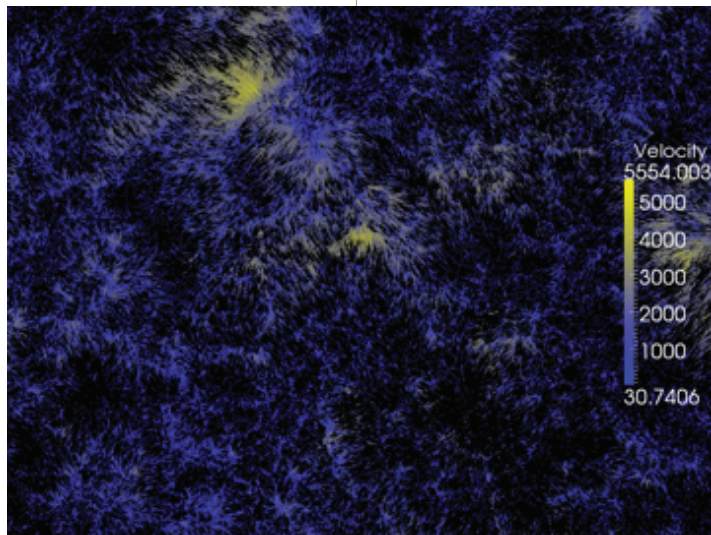
Funding Acknowledgments

LANL Directed Research and Development Program

The Roadrunner Universe Project: Baryon Acoustic Oscillations in the Intergalactic Medium

Adrian Pope, ISR-1/T-2/CCS-6; Katrin Heitmann, ISR-1; Salman Habib, Zarija Lukić, T-2; Patricia Fasel, CCS-3; David Daniel, CCS-1; Martin White, Jordan Carlson, University of California, Berkeley

Fig. 1. Dark matter halos from one of the large Roadrunner simulations, with 1/64 of the total $(750\text{Mpc}/h)^3$ volume displayed. The halos are shown as arrows, colored with respect to their velocity magnitude. This particular simulation was run with 64 billion particles, where each simulation particle has a mass of approximately one billion suns. The snapshot is taken at a redshift of $z=2.5$.



Over a decade ago, supernova observations first suggested that the expansion rate of the Universe was accelerating; this result has now been verified by a number of observations. The acceleration can either be driven by a mysterious source in the Einstein equations, the dark energy with a negative equation of state, or may hint at possible modifications of general relativity. Current datasets, while agreeing on the fact of the acceleration, only constrain the equation of state to about 10% [1]. At this level of accuracy, the acceleration could be explained by a cosmological constant – perfectly acceptable classically, but severely incompatible with quantum estimates by many orders of magnitude.

More accurate observations are needed to decide whether the equation of state is time variable (which would rule out a cosmological constant), or whether general relativity is valid on very

large cosmic scales. It turns out that the formation of structure can be an excellent probe of the evolutionary history of the Universe. Moreover, by tracking the growth of structure one can also test whether the predictions of general relativity are valid. For these reasons, cosmological probes based on structure formation dominate discussions of precision cosmology, targeting accuracy levels of 1% or better.

Analytical methods are of limited use at this level of accuracy, and one must resort to numerical simulations. To meet this challenge, a new hybrid petascale code has recently been developed by us at LANL, under the aegis of the Roadrunner Universe (RRU) project [2]. The code was tested and run on Roadrunner during the recent Open Science period when the full machine was made available to a select number of application codes.

The Open Science RRU run on Roadrunner consisted of nine ultra-large simulations to study the imprint of oscillations in the baryon-photon plasma in the early Universe, the Baryon Acoustic Oscillations (BAO). Due to BAO, a distinct but subtle signature is imprinted on the large-scale distribution of matter and has been seen in the spatial statistics of the distribution of galaxies [3], confirming one of the most important predictions of modern cosmology. BAO has now become one of the premier methods for determining cosmological distances, and hence, the expansion history of the Universe.

Traditional galaxy-based BAO surveys require a heavy investment in telescope time, especially as one goes to higher redshifts. Fortunately, tracers of the mass distribution, other than galaxies, exist. Neutral hydrogen in the intergalactic medium (IGM) furnishes one such example. At redshifts of $z \cong 2-3$, the gas making up the (IGM) is thought to be in photoionization equilibrium, resulting in a tight density-temperature relation, with the neutral hydrogen density proportional to a power of the baryon density [4]. Since pressure forces are subdominant, the neutral hydrogen density closely traces the total matter density on large scales. The neutral hydrogen density can be probed by obtaining spectra of distant, bright compact sources—the quasars—and studying the celebrated Lyman- α forest of absorption lines that map the neutral hydrogen along the line-of-sight to the quasar. The structure in quasar absorption thus traces, in a calculable way, slight fluctuations in the matter density of the universe back along the line-of-sight to the quasar, with most of the Lyman- α forest arising from overdensities of a few times the mean density.

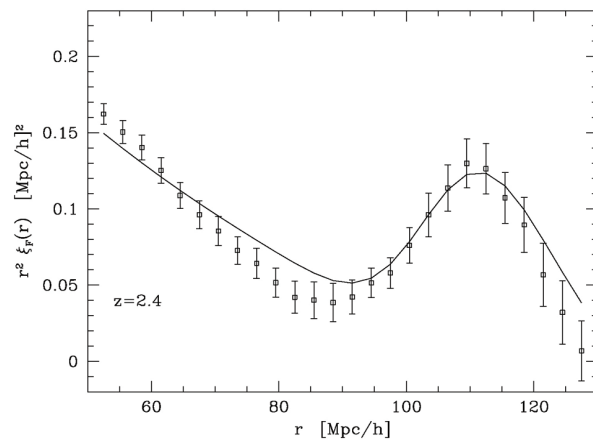


Fig. 2. The redshift space flux correlation function, ξ_F , as a function of the co-moving distance, r , measured in Mpc/h . The BAO feature is the bump centered around $110 h^{-1} \text{Mpc}$. The error bars are derived from a bootstrap analysis. The solid line is a Gaussian-smoothed linear theory result multiplied by a scale-independent bias to match the simulations.

The upcoming Baryon Oscillation Spectroscopic Survey (BOSS) [5] will provide an unprecedented number of quasar spectra for Lyman- α studies, motivating a major simulation effort at understanding the BAO imprint in the IGM. The set of Roadrunner simulations [6] are the first to simultaneously resolve structure down to the Jeans scale of the gas ($\cong 100 \text{ kpc}$), as well as properly capture the acoustic scale ($\cong 100 \text{ Mpc}$). Using the results of the simulations (the density and velocity fields), mock quasar spectra were constructed by running lines of sight from quasar sources to an observer. These spectra have properties close to those observed at $z \cong 2-3$. Because these mock spectra will be very useful in testing observational data pipelines, calibrating analysis tools, and in planning future projects, they have been made publicly available.

Given the quasar spectra derived from the simulations, one can compute the flux-flux correlation function, which is related to the underlying nonlinear, redshift-space, mass correlation function. The characteristic BAO bump signal in the flux correlation function, as measured from our simulations, is shown in Fig. 2. Although

our total simulation volume is large (nine $750h^{-1} \text{Mpc}$ boxes), covering effectively 1000 sq. deg. of sky, it is still only 10% of the area planned for BOSS. Thus the fractional errors in the flux correlation function as achieved by BOSS should be better than in our simulations by a factor of three.

Once detailed simulations are available, several important effects can be studied, including key sources of systematic errors and bias in the observations. We carried out preliminary investigations of an evolving mean flux, fluctuations in the photoionization rate, and HeII reionization, which generate extra power on the acoustic scale and reduce the contrast of the acoustic peak. Gravitational instability produces a well-defined pattern of higher-order correlations that is not obeyed by nongravitational contributions such as the above, allowing, in principle, a diagnostic of nongravitational physics in the forest. As an example, we demonstrated that the three-point cross-correlation function in models with HeII reionization has a different scale dependence than the three-point function in gravity-only simulations, regardless of the equation-of-state assumed in the latter.

As one part of the Roadrunner Universe project, future BAO simulations will be run with much larger boxes for an improved treatment of planned surveys. Additionally, the quasars will not be randomly placed in the simulation box, but will follow the appropriate statistical occupation distribution for being hosted by dark matter halos.

For more information contact Salman Habib at habib@lanl.gov.

-
- [1] E. Komatsu et al., *Astrophys. J. Supp.* **180**, 330 (2009).
 - [2] A. Pope et al., *Comp. Sci. Eng.* (to appear).
 - [3] D.J. Eisenstein et al., *Astrophys. J.* **633**, 560 (2005).
 - [4] A.A. Meiksin, *Rev. Mod. Phys.* **81**, 1405 (2009).
 - [5] D. Schlegel, M. White, D. Eisenstein, arXiv:0902.4680 [astro-ph.CO].
 - [6] M. White et al., *Astrophys. J.* (to appear); arXiv:0911.5341 [astro-ph.CO].

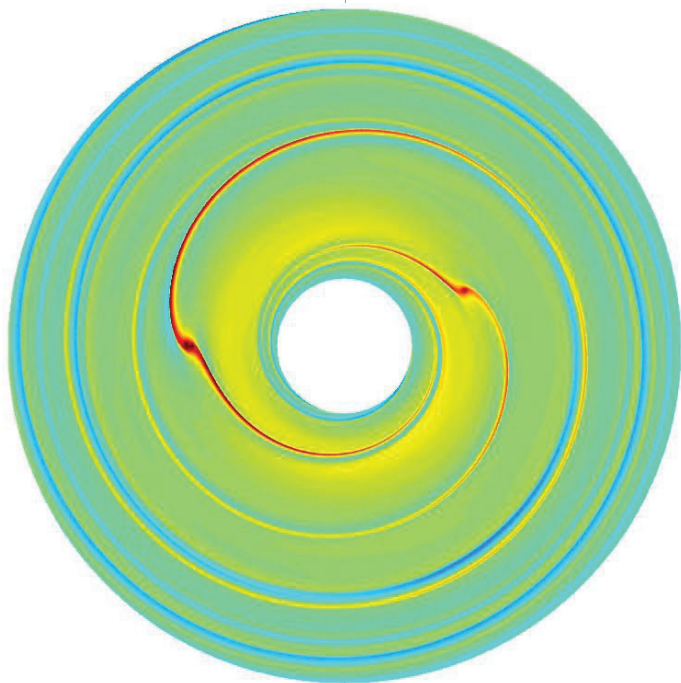
Funding Acknowledgments

- LANL Directed Research and Development Program
- NASA Theoretical Astrophysics Programs

Fast and Accurate Simulations of Proto-planet Migration in Disks

Shengtai Li, T-5; Hui Li, T-2

Fig. 1. Density plot for the tidal interaction between the disk and the embedded proplanets.



More than 400 extra-solar planets have been discovered since the first discovery in 1995. Planets are believed to form from the protoplanetary disks of gas and dust that are observed to orbit young stars. Once formed, planetary orbits may be modified as a result of gravitational tidal interactions with their nascent gaseous disk. This interaction can result in planetary migration (e.g., moving towards the parent star), directly threatening the survivability of protoplanets and altering their orbits. One of the primary challenges in this field is to understand the observed orbital properties of the more than 400 planets discovered that show marked differences from our own Solar system.

The interaction between the disk and planets (see Fig.1) is difficult to simulate. First, it requires an integration time up to thousands of orbits and millions of simulation time steps to obtain results that are insensitive to the initial conditions [1]. Second, it involves many physical processes that are represented either by partial differential equations with discontinuities for the disk or by ordinary differential

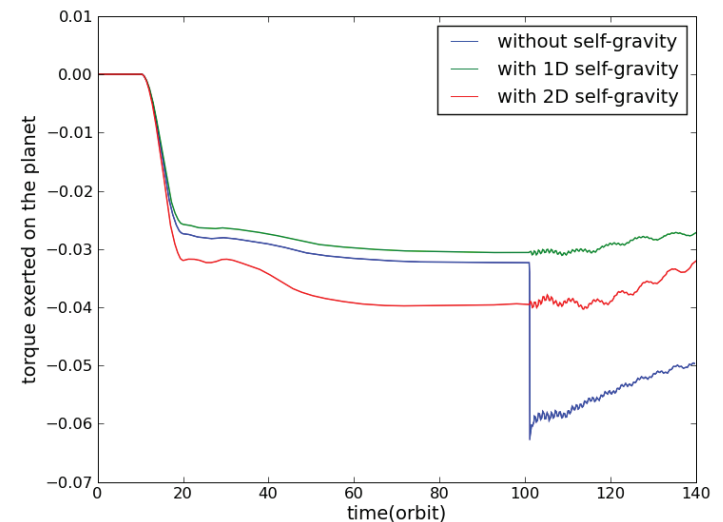


Figure 2. The impact of the disk self-gravity on torque exerted on the planet.

equations for the planets. The disk motion and the planet motion must be coupled tightly to obtain an accurate estimation of the torque exerted on the planet. Third, it requires a very high-resolution mesh, especially for the case of the high-mass planet, to resolve the near-planet region and the co-orbital region [2].

Thanks to the FARGO, fast advection in rotating gaseous objects, algorithm [3], we can simulate the interaction between the disk and planet using a fine-resolution grid for a long time. We have observed that FARGO is at least an order of magnitude faster than other standard solvers. We have modified the original FARGO so that it becomes more efficient and accurate [4].

The self-gravity of the disk is important in estimating the torque exerted on the planet. Without self-gravity, the coupled system is treated inconsistently because the planet is subject to the gravity of the disk, whereas the disk itself is not. We have implemented a full 2D self-gravity solver on a uniform disk grid [5]. This solver uses a mode cut-off strategy to reduce the computation and the

communication between different processors in parallel computing. The algorithm is sufficiently fast that the self-gravity solver costs less than 10% of the total computation cost in each run. Figure 2 shows the difference on torque estimation with and without disk self-gravity.

The planet motion requires a smaller time step than that of the disk motion. During one disk motion time step, the planet motion is calculated with high-order time integrations and controlled in a subcycling fashion by moving within a 0.05 grid spacing in each substep. The disk gravitational force on the planet, which is calculated whenever the disk and planet move to the same position in the azimuthal direction, is assumed to evolve linearly with time during these substeps. To enhance the resolution near the planet, we developed and implemented an embedded Lagrangian adaptive mesh refinement (AMR) for the FARGO-type solver. Unlike traditional AMR, our local refinement can move across several cells in the azimuthal direction in one time step. This combination of AMR with FARGO speeds up our computation by another order of magnitude.

The simulation is 2D but includes a prescription to approximate the effects of 3D when calculating the gravitational force from the disk material. This is accomplished by spreading the surface density of a given cell vertically in a Gaussian profile. With this pseudo-3D treatment migration rates from simulations with sufficient viscosity—dimensionless kinematic viscosity ($1\text{E}-6$) agrees well (within a few percent) with the 3D linear theory results of Tanaka et al. [6].

We have used our code to study the Type-I migration rate for the low-mass planet [1]. We found that Type-I planet migration can be halted in disks of sufficiently low turbulent viscosity. We have also applied the code to study the Type-III migration, and found that run-away migration is possible under certain circumstances [7].

For a relatively high-mass planet and disk with low or zero viscosity, we found that the shock-induced vortices can be generated in the disk (see Fig. 3). These vortices greatly affect the planet migration rate and even lead to gravitational instability.

The code is fully parallelized with a message-passing interface (MPI) and is verified using the published examples. More details about the algorithm and its numerical experiments can be found in [4].

For more information contact
Shengtai Li at sli@lanl.gov.

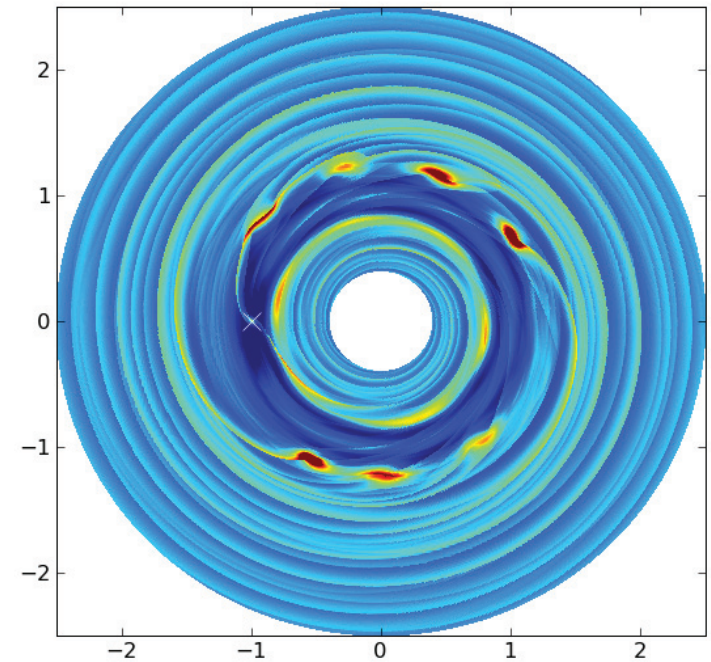


Fig. 3. The impact of viscosity on the planet migration rate.

- [1] H. Li et al., *Ap. J. Lett.* **690**, L52 (2009).
- [2] G. D'Angelo et al., *MNRAS* **358**, 316 (2005).
- [3] F. S. Masset, *A&A* **141**, 165 (2000).
- [4] S. Li, H. Li, "A Fast Parallel Simulation Code for Interaction between Proto-planetary Disk and Embedded Proto-planets," LA-UR 09-2968.
- [5] S. Li et al., *Ap. J.* **181**, 244 (2009).
- [6] H. Tanaka et al., *Ap. J.* **565**, 1257 (2002).
- [7] S. Li, H. Li, "Runaway Migration in a Massive Disk," LA-UR 09-06019; to be submitted to *Ap. J.*

Funding Acknowledgments

- Department of Energy
- LANL Directed Research and Development Program

Using Roadrunner to Model Shock Breakout in Supernovae and National Ignition Facility Experiments

Timothy M. Kelley, Todd Urbatsch, CCS-2; Aimee Hungerford, XTD-6; Gabriel Rockefeller, Chris L. Fryer, Paul J. Henning, Jeff Densmore, CCS-2; Barbara DeVolder, XCP-6

Thermal X-ray transport is an important physical component in phenomena, such as supernova explosions and inertial confinement fusion, where hot, radiating material is tightly and nonlinearly coupled to the radiation it emits. The time-implicit method for simulating thermal X-ray transport, Implicit Monte Carlo (IMC), approximates small-timescale absorption and reemission with a scattering process that is numerically stable but very computationally intensive [1].

The Jayenne IMC Project [2] was one of the early adapters for assessing the IBM Cell Broadband Engine (Cell BE) architecture for LANL's Roadrunner Project during the years 2005 to 2007. The Roadrunner supercomputer architecture increases a standard

Opteron cluster's floating point throughput by an order of magnitude with Cell BE processors. The early Cell BE-enabled version of the Jayenne Project stand-alone code, Milagro, adapted a limited set of capabilities to test the feasibility of scientific simulations on the Cell BE architecture. The Jayenne Project codes are general-capability IMC codes with multiple data-decomposition parallel schemes; multiple geometries of R, RZ, XYZ; multiple frequency

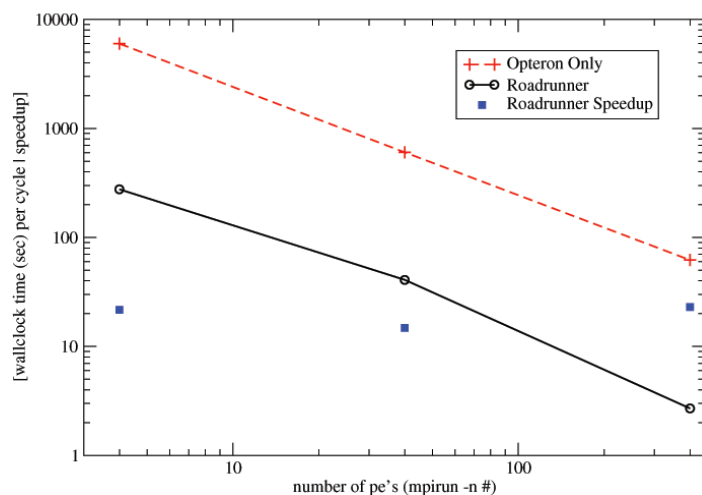
treatments, and a Random Walk method for speeding up the code in thick diffuse regimes, all of which can be run as the radiation-only stand-alone Milagro code or as the Wedgehog package for operator-split inclusion in a multiphysics application code. The first Milagro-Cell BE code was RZ-geometry, multigroup frequency treatment, replicated-geometry parallelism only, and absorption/reemission without physical scattering [3]. That first version showed speedups of up to eight on a radiation flow problem.

Since then, several additional Jayenne Project capabilities have been adapted to the Cell BE architecture, including domain-decomposition parallelism, physical scattering, the Random Walk speedup technique, and extension to the Wedgehog interface for radiation-hydrodynamics simulations via LANL's Cassio code, which models radiation, electrons, and ions independently in a nonequilibrium fashion.

The first comparison of this new Roadrunner-ready code with the standard version was on a simple radiation-only, steady-state, infinite homogeneous problem that, over time, remains at its initial temperature. Given that only the IMC particle transport utilizes the Cell BE chips, and that this problem is almost 100% particle transport, this problem represents an upper limit of the speedups with this version of Cassio/Wedgehog. The runtimes per cycle and speedups are both shown in Fig. 1, with speedups in the range of 15-22.

Our Roadrunner Open Science run focused on using the Cassio code to model the shock breakout in a supernova explosion. The engine of a supernova explosion is buried deep within the exploding star. Initially, the photons in the star are trapped in the outward flowing explosion. But as this explosion shock hits the surface of the star, the photon mean free path increases dramatically, allowing the photons to lead the explosion. This shock "breakout" is the first burst of photon radiation seen in the explosion and has now been observed in a number of supernovae. Until now, all models of this shock breakout were either limited to 1D simulations using flux-limited diffusion or equilibrium calculations that assumed the matter and radiation spectra were described by a single temperature.

Fig. 1. Roadrunner and Opteron-only runtimes per cycle and Roadrunner speedups for a simple radiation-only problem.



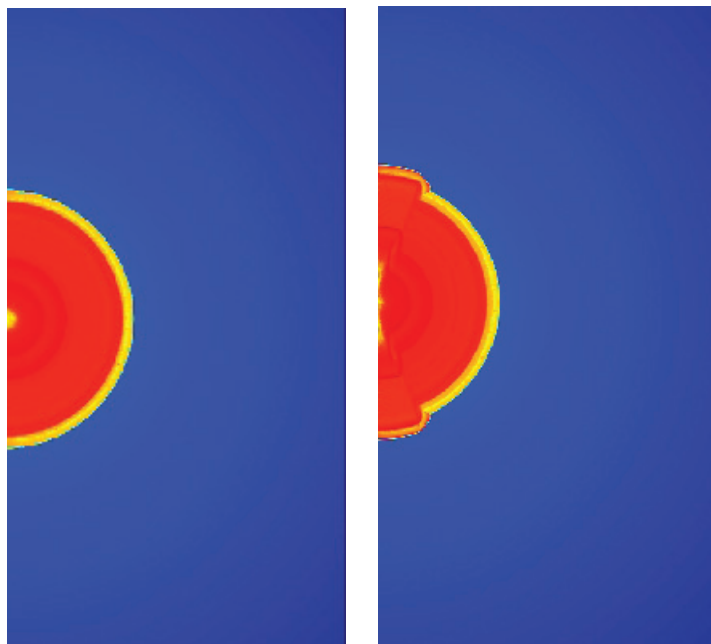


Fig. 2. Plot of the radial velocity structure (red denotes strong outward velocity, blue is zero velocity) of the shock breakout from two different supernova explosions: a symmetric explosion (left) and an asymmetric bipolar explosion based on the asymmetries studied in Hungerford et al. (2003).

Our calculations mark the first nonequilibrium radiation-hydrodynamics calculations of a supernova explosion in 2D, and the first-ever radiation hydrodynamics calculations of shock breakout with a higher-order transport scheme.

In Fig. 2, two of our calculations just after shock breakout show the velocity structure (colored by the magnitude of the radial velocity) for a symmetric and an asymmetric explosion using the bipolar explosion paradigm discussed in Hungerford et al. in an effort to match the asymmetries observed in supernova 1987A [4]. There is growing evidence that many core-collapse supernovae have such asymmetries. But without detailed spectral calculations, it will be difficult to determine the exact nature of the asymmetry. Our simulations mark the beginning of such precision comparisons of theory and observations. The NASA Swift satellite has begun an

active observing program to observe shock breakout, and we are working closely with this team to couple theoretical models to the observations. This effort is a key aspect of the SciDAC (DOE's Scientific Discovery through Advanced Computing program) priorities identified in nuclear astrophysics and will be part of the upcoming SciDAC science competition.

This capability is also being used to simulate experiments on the National Ignition Facility (NIF). For this calculation on 32 processors, the standard version of the code took 37 hours to run. Running the same input deck on the Roadrunner version took about 5 hours, yielding an overall speedup of seven. Figure 3 shows the running wallclock time for each code version.

As we look beyond Roadrunner to the next generation of heterogeneous architectures, we will look to speed up a greater fraction of the overall computation and develop new methods that were maybe too compute-intensive to consider before Roadrunner, but that now might make sense.

**For more information contact
Tim Kelley at tkelley@lanl.gov.**

Roadrunner Speedups of 3-7 for a NIF Simulation

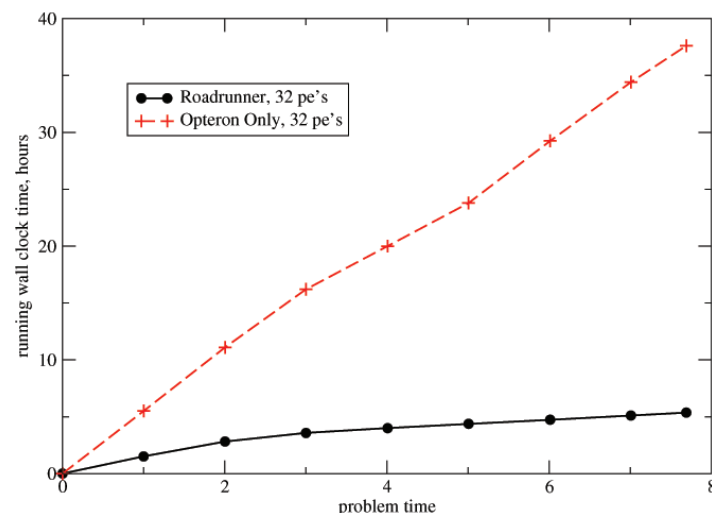


Figure 3. Cumulative wallclock times over the duration of a NIF experiment simulation for 32 processors on both Roadrunner and Opteron only. Cumulative Roadrunner speedups are about three early in the problem, and seven at the end of the problem.

- [1] J.A. Fleck, Jr., J.D. Cummings, "An implicit Monte Carlo scheme for calculating time and frequency dependent nonlinear radiation transport", *Journal of Computational Physics*, **8**, 313 (1971).
- [2] T.J. Urbatsch, T.M. Evans, "Milagro Version 2, An Implicit Monte Carlo Code for Thermal Radiative Transfer: Capabilities, Development, and Usage," Los Alamos National Laboratory report LA-14195-MS (Jan. 2005).
- [3] T.M. Kelley and P.J. Henning, "Adaptation of Milagro Implicit Monte Carlo to the Roadrunner Hybrid Architecture," in *ADTSC Science Highlights 2008*, Los Alamos National Laboratory, LA-UR-08-1960.
- [4] A.L. Hungerford, C.L. Fryer, M.S. Warren, *Ap. J.*, **594**, 390 (2003).

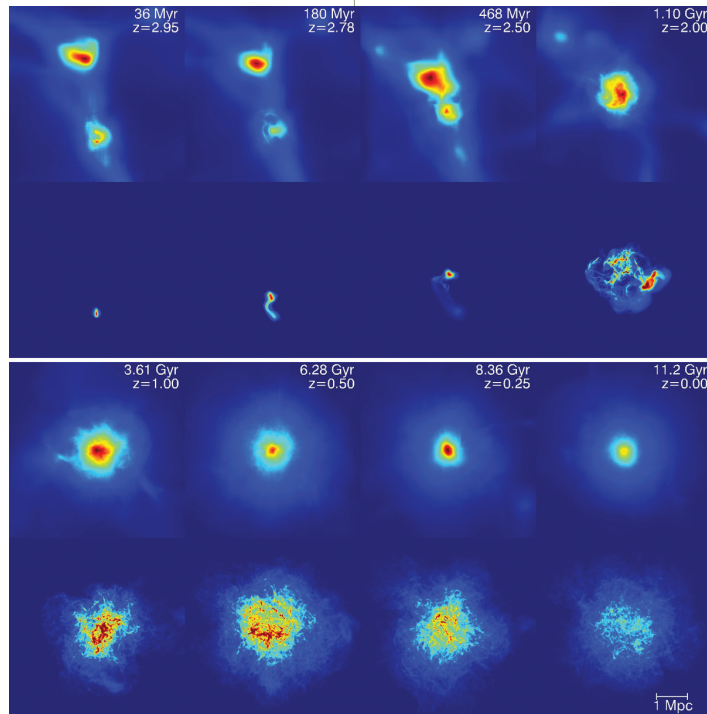
Funding Acknowledgments

- LANL Directed Research and Development Program
- LANL Institutional Computing Program

Origin of Magnetic Fields in Galaxy Clusters

Hao Xu, Hui Li, T-2

Fig. 1. Snapshots of the projected baryon density (upper rows) and magnetic energy density (lower rows) for different epochs of cluster formation. Each image is 5.71 Mpc (comoving) on a side. Cluster mergers continuously extend the distribution of the magnetic fields and amplify them.



Galaxy clusters are the largest gravitationally bound objects in the Universe. Each cluster contains as many as thousands of galaxies and large amounts of intergalactic gas known as the intracluster medium (ICM). The detection of large-scale, diffuse radio emissions, called radio halos, and relics in galaxy clusters, as well as other observations, show that the ICM is permeated with cluster-wide micro Gauss (μG) magnetic fields [1]. These magnetic fields play a significant role in determining the structure of clusters through processes such as heat transport, which consequently affect

the applicability of clusters as sensitive probes for cosmological parameters [2].

The origin of cluster magnetic fields is still not clear. Two possibilities have been considered promising and have received the most attention: 1) magnetic fields are initially from the outflows of normal or active galaxies, and such fields can be further amplified by cluster turbulence, and 2) very weak proto-galactic seed fields, such as from the Biermann battery effect, are amplified by dynamo processes in clusters.

In recent work, we have made significant progress in addressing this origin problem by showing

that the observed magnetic fields can be obtained by amplifications of magnetic fields that are injected by an active galactic nucleus (AGN) using large-scale cosmological simulations [3]. We used the newly developed ENZO+MHD code, which is a grid-based cosmological Magneto Hydrodynamics (MHD) + N-body code with adaptive mesh refinement (AMR) [4]. This code was developed jointly by the University of California, San Diego and LANL. It uses the AMR algorithm to improve spatial and temporal resolution in regions of interest, such as gravitationally collapsing objects. In this research, we performed a detailed self-consistent simulation of galaxy cluster formation without magnetic fields from redshift $z = 30$ to $z = 3$. Then we restarted the simulation with the magnetic fields injected into the largest halo at that time, using a magnetic tower model [5] to mimic the magnetic fields output from an AGN, and ran the simulation to $z = 0$. The magnetic field injection lasts for 36 Myr and the injected magnetic energy is $\sim 2 \times 10^{60}$ ergs.

The formation of clusters and the evolution of magnetic fields are shown in Fig. 1 as snapshots of projected gas density and magnetic energy density. The injected local magnetic fields are distributed throughout the entire cluster when the cluster is formed by hierarchy mergers. After the injection, the expansion of the injected magnetic fields first forms the density cavities, reminiscent of the jet-lobe structure of an AGN outburst. As the evolution progresses, magnetic fields, which follow the gas motion, are being sheared, twisted, and spread throughout the whole cluster. Judging by the images, this volume-filling process is quite efficient. By $z = 0.5$, magnetic fields are already well mixed with the ICM and are distributed quasi-uniformly throughout the whole cluster, except that some high magnetic field regions obviously from compression by merger shocks. At the time when the simulation ends, while the cluster is relaxed, the magnetic fields are quite uniformly distributed.

During the course of the magnetic fields spreading throughout the whole cluster, the total magnetic energy increases significantly. The magnetic energy inside the cluster increases to more than 10^{61} ergs with a 25 times gain (see the top panel of Fig. 2). The spherical averaged radial profiles (bottom panel of Fig. 2) of magnetic field

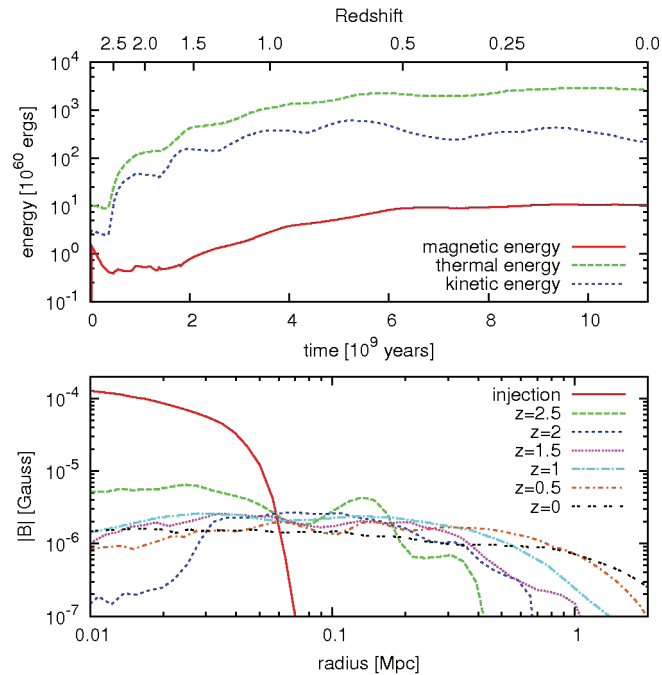


Fig. 2. *Top panel: Temporal evolution of different components of energy inside the virial radius of the cluster. Bottom panel: The spherically averaged radial profile of magnetic field strength. The radius is measured in the proper frame.*

strength at different epochs of the cluster formation show the turbulent diffusion of magnetic fields throughout the cluster, yet maintain their strength via the dynamo process. The strength of magnetic fields at $z = 0$ is a few μG at the cluster center and drops slowly with increasing radius to sub- μG at the cluster edge. This radial profile is consistent with the results from observations.

We also computed the synthetic Faraday rotation measurement (FRM) of the cluster, which is the most-used observation for determining the structure of magnetic fields in cosmological objects. The typical value of FRM (Fig. 3) in our simulation is $\pm 200 \text{ rad m}^{-2}$, with high values concentrated in the cluster core region. The FRM map not only shows the small-scale variations that are reminiscent of MHD turbulence of the ICM, but also

displays long, narrow filaments. The FRM magnitudes and spatial distributions from our simulation are quite consistent with observations of clusters.

Since AGNs are commonly observed in galaxy clusters, one important implication of our studies is that the magnetic fields from AGNs alone are perhaps enough to provide the initial magnetic fields in the clusters, and the ICM turbulence will spread and amplify the AGN fields to the observed magnetic fields in clusters via a dynamo process. Future radio and X-ray observations, especially the Extend Very Large Array, which is under construction, will provide an unprecedented level of detail of radio sources and magnetic fields in the ICM. Such theoretical and simulation studies will be timely in understanding the observations.

For more information contact Hao Xu at hao_xu@lanl.gov.

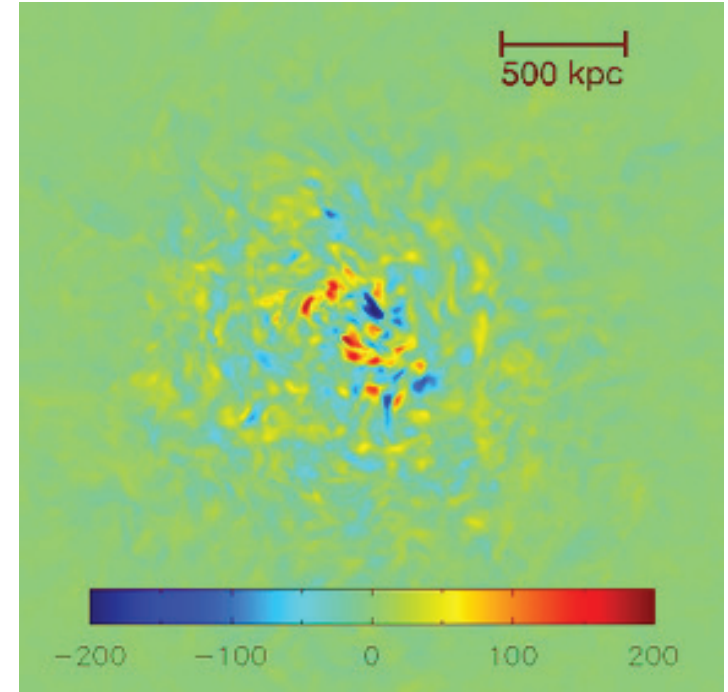


Fig. 3. *Faraday rotation measurement of the ICM of the simulated cluster at $z = 0$. The distribution of FRM is consistent with the observed FRM distributions from real clusters.*

- [1] C.L. Carilli, G.B. Taylor, *Annu. Rev. Astron. Astrophys.* **40**, 319 (2002).
- [2] G.M. Voit, *Rev. Mod. Phys.* **77**, 207 (2005).
- [3] H. Xu et al., *Astrophys. J. Lett.* **698**, 14 (2009).
- [4] D.C. Collins et al., *Ap J Supplement*, in press (2010).
- [5] H. Li et al., *Astrophys. J.* **643**, 92 (2006).

Funding Acknowledgments

- LANL Directed Research and Development Program
- LANL Institute of Geophysics and Planetary Physics (IGPP)
- DOE Fusion Energy Sciences Program
- LANL Institutional Computing Program

Atomic, Nuclear, and High Energy Physics

The first article in this section describes an approach to study infrequent-event systems, such as the diffusion of atoms on surfaces. Next is an article on the simulation of X-ray diffraction patterns from large-scale molecular dynamics simulations. The third article describes a combined quantum mechanics-molecular dynamics technique with very efficient performance. This approach has applications in the study of assemblies of organic molecules. The last article describes work on the ionization of atom by electron impact, a topic of importance in astrophysics as well as the nuclear weapons programs.

Atomic, Nuclear, and High Energy Physics

Ab Initio Parallel Replica Molecular Dynamics Method

Elena Jakubikova, Danny Perez, Arthur F. Voter, Richard L. Martin, Enrique R. Batista, T-1

Dynamical behavior of infrequent-event systems is characterized by a trajectory that resides in a potential basin for many vibrational periods before undergoing a transition to a new basin. Modeling dynamics of the infrequent-event systems is very challenging, since very large time scales need to be considered in order to capture their relevant behavior. One of the ways to tackle this challenge is through the use of accelerated molecular dynamics methods, such as parallel replica molecular dynamics (PRMD) [1].

The PRMD method parallelizes the molecular dynamics (MD) simulation in the time domain. The only assumption made is that the infrequent events studied obey the first-order kinetics. The PRMD method is illustrated in Fig. 1. The system studied is first prepared in an initial state, which is then replicated into M copies sent to M processors. All copies (or replicas of the system) are then dephased in order to eliminate correlations between them. After the dephasing stage, each processor carries out an independent MD simulation, exploring the phase space in the original basin

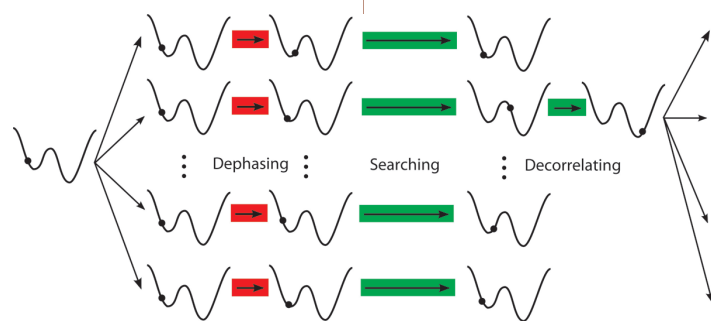
M times faster than a single MD trajectory would. Once the transition is detected, the simulation clock is advanced by the sum of the accumulated MD time in each replica. The trajectory that made the transition is allowed to continue for an additional amount of

time called decorrelation time, during which recrossing or other correlated events may occur. The simulation clock is then advanced by decorrelation time, the system in the new state is again replicated into M copies, and the whole process is restarted.

To obtain forces acting on atoms in each step of an MD simulation, current implementations of the PRMD method rely on the use of empirical potentials fitted to empirical data or electronic structure calculations. Accurate empirical potentials are hard to develop and they must be redetermined if one changes an atom type. Therefore, ideally, one would like to use forces calculated from ab initio electronic structure calculations. Of course, the calculation of ab initio forces is much more expensive, but with the help of currently available massively parallel supercomputers there are many problems that can be tackled this way. We have written an implementation of PRMD that obtains forces on the fly at each simulation step from ab initio packages. This implementation was done by utilizing a glue code, written in Perl, that ties together the PRMD code with an electronic structure code (see Fig. 2). At each step of an MD simulation, the PRMD code passes the current positions of the simulated system to the Perl script, which processes them and sends them to an electronic structure code. After the electronic structure code calculates the current forces, they are passed via the same script back to the PRMD code. Our implementation of the ab initio PRMD method contains two levels of parallelism: first, it parallelizes the molecular dynamics simulation in the time domain by the use of multiple replicas; second, the energies and forces at every step are calculated by the use of parallel electronic structure codes. Ab initio PRMD has a very favorable scaling with respect to the number of processors and, depending on the type of the system, hundreds (and even thousands) of processors can be used to perform the simulations.

The ab initio PRMD was implemented on the Ranger supercomputer at the Texas Advanced Computer Center and applied to the study of diffusion dynamics of helium in hexagonal ice. The Vienna Ab-initio Package Simulation (VASP) program [2, 3] was

Fig. 1. Schematic of the PRMD method. Total simulation time is a sum of MD time spent searching or decorrelating (green arrows). Dephasing time (red arrows) does not advance the simulation clock.



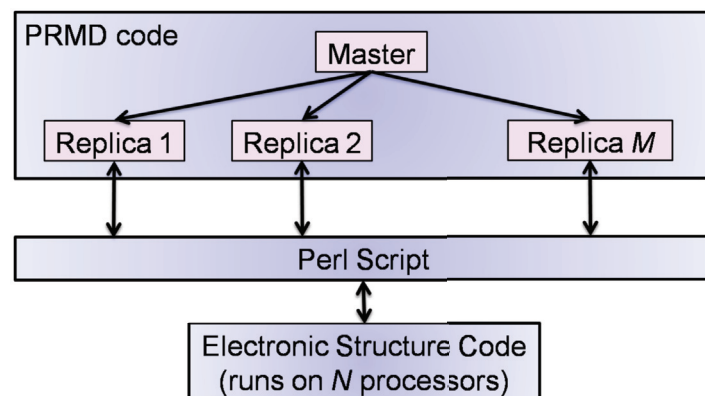


Fig. 2. *Ab initio* PRMD implementation.

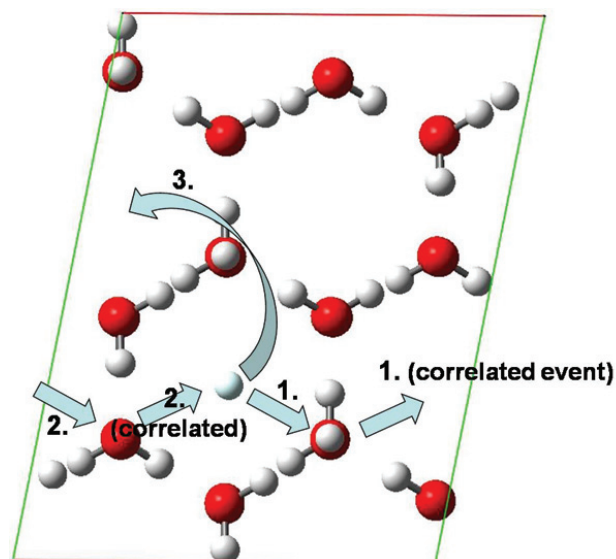


Fig. 3. Helium in hexagonal ice. Blue arrows show observed transitions.

used to obtain forces at every step of the MD simulation. One hundred replicas were used to parallelize the MD simulation in the time domain, with each electronic structure calculation parallelized on 16 processors (16 processors equal one node on Ranger),

leading to the total use of 101 nodes (one for each replica and a master node) corresponding to the use of 1616 processors total. The number of replicas was reduced to 20 in the later stages of the simulation. The preliminary results of the MD simulation are shown in Fig. 3. The simulation was done at 263K, with a 0.5 fs time step. One-picosecond blocks were used for dephasing and decorrelating. The total MD time needed to observe the three transitions depicted in Fig. 3 was 340 ps. With each time step lasting approximately 140 s of wall time, it would take 3.02 yr to reach this time with a standard MD simulation—however, it took only 42 days with our implementation of the ab initio PRMD methodology.

In conclusion, we have implemented the ab initio PRMD method, which is capable of simulating infrequent event processes. We have tested this implementation by studying diffusion of helium in hexagonal ice and shown that the ab initio PRMD is capable of reaching time scales several order of magnitudes longer than standard ab initio MD methods while maintaining full atomistic description of the simulated system.

For more information contact Enrique R. Batista at erb@lanl.gov.

- [1] A.F. Voter et al., *Ann. Rev. Mater. Res.* **32**, 321 (2002).
- [2] G. Kresse, J. Furthmüller, *Comput. Mater. Sci.* **6**, 15 (1996).
- [3] G. Kresse, J. Furthmüller, *Phys. Rev. B* **54**, 11169 (1996).

Funding Acknowledgments

- LANL Directed Research and Development Program
- National Science Foundation (NSF)

Simulation of X-ray Diffraction Patterns

John L. Barber, T-1

For more than a century the experimental technique of X-ray diffraction has been an invaluable tool for both diagnostics and imaging of diverse types of samples. However, there is often a need to generate an X-ray diffraction pattern corresponding to a given hypothetical experimental setup without directly performing the actual experiment. Such a need arises, for example, in the design of experiments or experimental facilities, or in the analysis or validation of large-scale molecular dynamics (MD) simulations. In this highlight, I give a brief overview of some recent efforts in the computer simulation of X-ray diffraction patterns produced from arbitrary samples and with a wide range of sources and detectors.

Suppose we have a collection of N atoms, with coordinates $\{\vec{x}_j\}$ (in practice this set of atomic coordinates is most often obtained from an MD simulation), as well as the schematic experimental setup shown in Fig. 1. If we assume for the moment that the light source is a monochromatic, unpolarized, coherent beam, it is straightforward to show that the net scattered intensity at a position \vec{y} on the detector is given by

$$I(\vec{y}) = I_o \left(\frac{r_e}{|\vec{y}|} \right)^2 \left(1 - (\hat{\epsilon} \cdot \vec{y}) / |\vec{y}| \right) \cos(2\theta) \left| \sum_{j=1}^N Z_j e^{-i\vec{q} \cdot \vec{x}_j} \right|^2 \quad (1)$$

where I_o is the intensity of the source, r_e is the classical radius of the electron, $\hat{\epsilon}$ is the polarization vector of the source, 2θ is the deflection angle of the beam, Z_j is the atomic number of the j^{th} atom, and the scattering vector $\vec{q} = \vec{k}_j - \vec{k}_i$ is the difference between the final and initial wave vectors of the scattered light.

The various terms in Eq. (1) are straightforward to evaluate, except for the sum over particles, which can only be evaluated by direct "brute-force" means for systems consisting of less than 10,000 particles. For each point \vec{y} on the detector, there corresponds a distinct scattering vector \vec{q} . Assuming a 300×300 desired resolution in the diffraction pattern, and 10^9 atoms in the sample (a common number in modern MD simulations), this sum amounts to a prohibitive 10^{14} separate terms. Note that although it has the appearance of a Fourier transform, this sum cannot be directly evaluated via fast Fourier transform techniques, since the atomic positions are in general arbitrarily spaced, and thus do not fall precisely on a grid. This problem can be sidestepped, however, via the observation that the sum

$$S(\vec{q}) \equiv \sum_{j=1}^N Z_j e^{-i\vec{q} \cdot \vec{x}_j} \quad (2)$$

may be rewritten as

$$S(\vec{q}) = \int d\vec{x} e^{-i\vec{q} \cdot \vec{x}} \rho(\vec{x})$$

where

$$\rho(\vec{x}) = \sum_{j=1}^N Z_j \delta(\vec{x} - \vec{x}_j) \quad (3)$$

is the number density of electrons in the sample. Equation (3) may be cast onto a regular grid by approximating each δ function via a suitable "quasi- δ " envelope function. In [1], for example, a Gaussian envelope is recommended, which has the advantage of speed and ease of implementation, but also has the drawbacks that it entails high memory usage (a concern in situations with larger data sets) as well as unavoidable approximations. In this work, a periodic sinc function envelope is used instead, which can be shown to both minimize memory usage, and allow an *exact* evaluation of the sum in Equation (2), at the cost of moderately higher computational time. Some example diffraction patterns generated by these methods from a 5.2-million atom MD simulation of a shock passing through copper and illuminated by a coherent 100 KeV beam are shown in

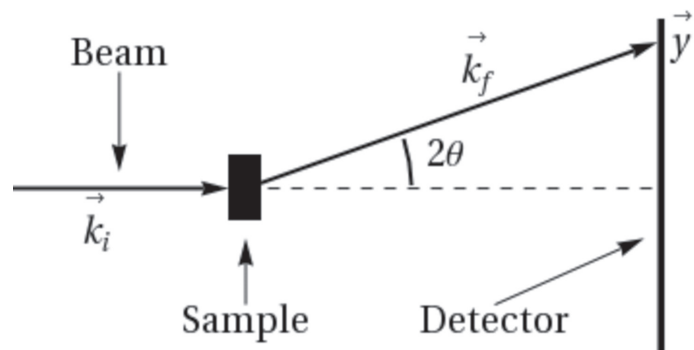


Fig. 1. A schematic diagram of the experimental setup simulated in this work.

Fig. 2. Each diffraction pattern took approximately 10 minutes to generate on a single-processor 3 GHz workstation. Data sets from MD simulations containing up to 40 million particles have been handled without difficulty.

There are a large number of variations of this problem (not described here due to limitations on length) that allow the consideration of, for example, arbitrary polarization spectra (the beam in Fig. 2 is assumed to be unpolarized), arbitrary broadband source spectra (i.e., Laue diffraction), virtual powderization of the sample, and point light source. See the "MPDH CXDI Analysis" page on the MaRIE wiki (<http://marie-sp.lanl.gov/wikimarie>) for more details.

For more information contact John L. Barber at jlbarber@lanl.gov.

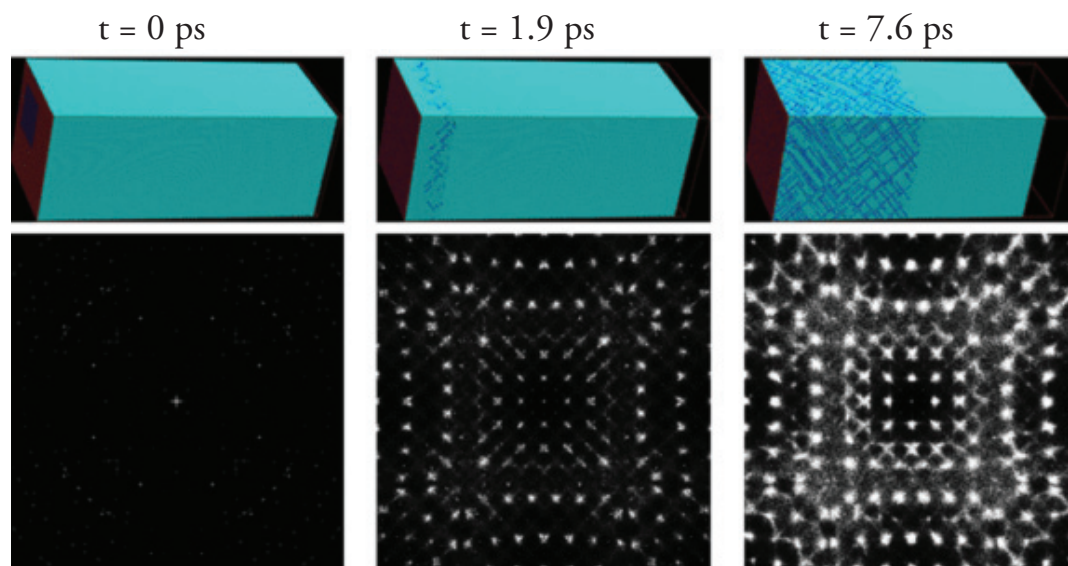


Fig. 2. Three snapshots from a 30 nm x 30 nm x 70 nm, 5.2-million atom MD simulation of a shock passing through FCC copper, along with the simulated X-ray diffraction patterns from each.

[1] G. Kimminau et al., *J. Phys. Condens. Matter* **20**, 505203 (2008).

Funding Acknowledgments

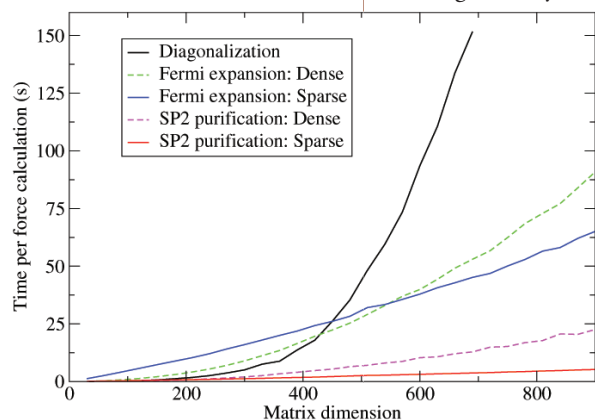
LANL MaRIE Capability Development

Linear Scaling Quantum Molecular Dynamics for Organic Molecules

Edward Sanville, Nicolas Bock, Matt Challacombe, Anders M. N. Niklasson, Marc Cawkwell, T-1

**For more information contact
Marc Cawkwell at
cawkwell@lanl.gov.**

Fig. 1. CPU time per MD time step as a function of matrix dimension using diagonalization of the Hamiltonian matrix, dense and sparse matrix implementations of SP2 purification, and the recursive expansion of the Fermi operator to calculate the density matrix.



Interatomic bonding in organic molecular crystals involves short-range, angularly dependent, intramolecular covalent bonds, medium-range van der Waals bonding, and long-range electrostatic interactions that arise from atomic partial charges. A predictive model for the structure and chemistry (i.e., the making and breaking of covalent bonds) of molecular crystals must describe all of these interactions, and their interdependencies, with physical rigor while ensuring computational tractability.

The selection of a mathematical model for interatomic bonding in materials typically involves a compromise between accuracy and computational speed. For example, ab initio quantum mechanical methods based on the self-consistent field or density functional theories are extremely accurate but computationally very expensive. Furthermore, the computational cost of these methods usually scales with the cube of the number of atoms, $O(N^3)$, making large-scale, long duration molecular dynamics (MD) simulations impossible. Parallel implementations of empirical interatomic potentials facilitate long duration MD simulations with many millions of atoms [1], but the ability of empirical potentials to describe a given material with high fidelity is often questionable. Furthermore, empirical

potentials tend to exhibit poor transferability to structures that differ from those employed in their parameterization.

In response to the limitations inherent to these two extremes, we have developed an $O(N)$ formalism and a supporting code that aims to strike a compromise between computational speed and physical accuracy. While at first it may seem strange to wish to compromise on the physical accuracy of an interatomic potential, a formalism that

is limited severely in the number of atoms that can be employed in a simulation will inevitably lead to errors associated with finite size effects. In many situations, errors originating from finite size effects can be more severe than those originating from approximations made during the construction of the interatomic potential.

We have generated a set of interatomic potentials for hydrocarbons based on the self-consistent charge transfer tight-binding (SCC-TB) approximation to density functional theory [2,3]. The SCC-TB formalism captures naturally the formation of covalent bonds and the transfer of charge between species of different electronegativity. The SCC-TB method is explicitly quantum mechanical since we construct and solve a time-independent Schrödinger equation for a system of interacting, atom-centered valence electrons, but the elements of the Hamiltonian matrix are parameterized rather than computed exactly. Semi-empirical potentials such as those based on tight-binding tend to show better transferability than empirical potentials owing to their more rigorous theoretical underpinning. They are also orders of magnitude faster than ab initio methods since many difficult-to-calculate terms are approximated and parameterized [4].

Our SCC-TB formalism is expressed in a minimal basis of an orthogonal set of real free-atom-like orbitals. The non-spin polarized SCC-TB Hamiltonian, H , is a sum of a Slater-Koster Hamiltonian that describes the formation of covalent bonds, H^0 , and the electrostatic potential generated by partial charges, q_i , on each atom,

$$H_{i\alpha,j\beta} = H_{i\alpha,j\beta}^0 + \frac{1}{2} \sum_k \gamma_{ik} q_k \delta_{ij} \delta_{\alpha\beta}$$

where i and j label atoms, α and β label orbitals, s, p_x, p_y, p_z , etc., γ_{ii} is the Hubbard U for species i , and $\gamma_{ij} = 1/R_{ij}$, where R_{ij} is the scalar distance between atoms i and j . The total energy is then

$$E_{\text{total}} = 2\text{tr}[\rho H] - \frac{1}{2} \sum_{i,j} \gamma_{ij} q_i q_j + E_{\text{pair}}$$

where ρ is the density matrix calculated from H and the electronic occupation. The SCC-TB equations must be solved self-consistently since the partial charges on which the Hamiltonian depends are obtained from the density matrix, i.e., $q_i = 2 \sum_{\alpha \in i} [\rho_{i\alpha,i\alpha} - \rho_{i\alpha,i\alpha}^0]$, where ρ^0 is the density matrix for isolated atoms. The last term in the

expression for the total energy, E_{pair} , is a sum of pair potentials that are used to provide short-range repulsion and $-1/R^6$ van der Waals interactions at long range.

In contrast to ab initio electronic structure methods, the construction of the Hamiltonian in SCC-TB is computationally very cheap. Instead, the bottleneck in SCC-TB calculations is the $O(N^3)$ computation of the density matrix. In the Los Alamos Transferable Tight-binding for Energetics (LATTE) code, we have implemented a number of methods for rapid, $O(N)$ calculations of the density matrix [5].

The most straightforward method for calculating the density matrix involves the diagonalization of the Hamiltonian [6]. While this method is simple to implement and enables a finite electron temperature to be easily incorporated, it is unsuited to calculations including large numbers of atoms because of $O(N^3)$ scaling (see Fig. 1). We find that Niklasson's second-order spectral projection (SP2) density matrix purification method is by far the most efficient method for the computation of the density matrix for large systems [7]. The starting point for SP2 purification is a remapping of all eigenvalues of the Hamiltonian into the interval $[0,1]$. This matrix is purified such that the eigenvalues of occupied states are projected toward 1 and those of unoccupied states toward zero until an idempotent density matrix is obtained. The purification operations require only one $O(N^3)$ matrix-matrix multiplication per iteration. However, for molecular systems we take advantage of the linear increase in the number of non-zero matrix elements with the number of molecules to achieve $O(N)$ matrix-matrix multiplications. Figure 1 illustrates clearly the tremendous speed-up that density matrix purification offers with respect to diagonalization for large systems even in dense matrix form, while a sparse matrix implementation yields $O(N)$ performance.

The density matrix obtained via purification corresponds to zero electronic temperature. This may lead to numerical instabilities during a MD simulation if eigenvalues cross the chemical potential. It is usually desirable to smooth these electronic transitions through the introduction of a finite electronic temperature. We find that the

recursive expansion of the Fermi operator method described in [8] allows for the calculation of a finite temperature density matrix with $O(N)$ cost for sparse molecular systems. It is clear from Fig. 1 that the recursive expansion of the Fermi operator is somewhere between diagonalization and purification in terms of speed, and that linear scaling is readily achievable.

Like SP2 purification, the Fermi operator expansion algorithm depends only on $O(N^3)$ dense matrix-matrix multiplications. This leaves LATTE with a very well defined computational bottleneck upon which we can concentrate our efforts. Thus, a key factor in the performance of LATTE is the question of how fast one can multiply square matrices. In addition to exploiting the sparsity of density matrices for molecular materials, we have found that the utilization of hybrid computational architectures is a promising route toward maximizing the speed of matrix-matrix multiplications. In Fig. 2 we show computational time as a function of matrix dimension for the double-precision general matrix multiply (DGEMM) algorithm on the Opteron central processing units (CPU) and IBM Cell Broadband Engine (Cell BE) chips on LANL's Roadrunner supercomputer. It is evident that the use of special computational architectures for time-consuming computational operations offers very significant performance gains. In parallel to this work on Roadrunner, we are actively pursuing the use of commodity graphics cards as computational accelerators such that huge performance boosts can be obtained even on off-the-shelf desktop workstations. Of particular note in this regard is the stability of LATTE's algorithms when executed in single-precision arithmetic.

A judicious selection of a physical model for interatomic

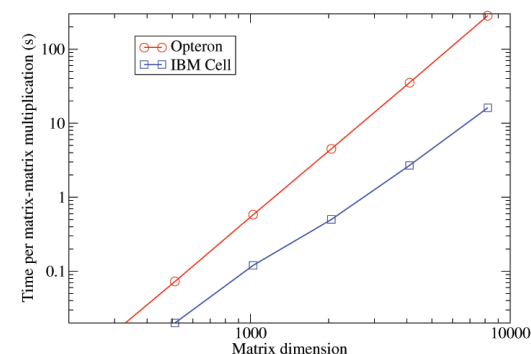


Fig 2. Time to compute a dense matrix-matrix multiplication using the DGEMM algorithm as a function of matrix dimension on the Opteron and Cell BE chips on Roadrunner.

bonding in molecular materials coupled to advanced algorithms, developments in applied mathematics, and novel computational architectures enabled us to design and run very large-scale quantum MD simulations. These simulations are currently being applied to the study of the chemistry under extreme conditions.

- [1] M.J. Cawkwell et al., *Phys. Rev. B* **78**, 014107 (2008).
- [2] M. Elstner et al., *Phys. Rev. B* **58**, 7260 (1998).
- [3] M.W. Finnis, *Interatomic Forces in Condensed Matter*, Oxford University Press (2003).
- [4] M. J. Cawkwell et al., *Phys. Rev. B* **73**, 064104 (2006).
- [5] E.J. Sanville et al., LATTE, LA-CC 10-004, Los Alamos National Laboratory.
- [6] A.P. Sutton et al., *Philos. Mag. A* **81**, 1833 (2001).
- [7] A.M.N. Niklasson *Phys. Rev. Lett. B*, **66**, 155115(2002).
- [8] A.M.N. Niklasson, *J. Chem. Phys.* **129**, 244107 (2008).

Funding Acknowledgments

- LANL Directed Research and Development Program
- DOE, NNSA Advanced Simulation and Computing (ASC) Program

Deep Interference Minima in Differential Cross Sections for Electron-impact Ionization

James Colgan, T-1

A recent paper published by James Colgan and coauthors as a “Fast-Track Communication” [1] has provided some insight into a long-standing puzzle in atomic collision physics.

The paper investigates an unusual deep minimum found in the triple differential cross sections (TDCS) resulting from the electron-impact ionization of helium for noncoplanar electron geometries. Such TDCS minima, which are effectively zeros in the cross section, are usually the result of selection rules (most often found in photoionization studies) or resulting from strong electron-electron correlation effects, where the two outgoing electrons repel each other at a given geometry. The minimum explored here, and as shown in Fig. 1, was discovered experimentally more than 15 years ago [2,3], and its existence has been difficult to explain, as no obvious selection rules or correlation effects seemed to play a role in the phenomena. The investigations by Colgan et al. [1] used an accurate nonperturbative time-dependent close-coupling (TDCC) approach to compute the triple differential cross sections for the ionization of helium at several incident energies, and for the various geometries of the experiment. The experimental data was reproduced extremely well by the TDCC calculations, as shown in Fig. 1. Figure 1 shows the TDCS for various incident electron energies, for given gun angles ψ as indicated (where the gun angle is the angle made between the incident electron beam and the plane of detection of the two outgoing electrons). The TDCS is presented as a function of the angle ξ , where 2ξ is the angle between the outgoing electrons.

Further analysis showed that the minima were due to strong destructive interference between the many partial waves (approximately 10 total angular momentum L terms are included for the calculations shown in Fig. 1) contributing to the TDCS. The strongly positive contributions from the direct terms in the partial wave sum are almost exactly cancelled by the negative contributions arising from the interference terms. Such destructive interference has also been found to contribute to other unusual TDCS phenomena, such as the rapidly changing TDCS when the bond length is varied in studies of double photoionization of molecular hydrogen [4].

Colgan et al. [1] also found that the deep minimum is present in the electron-impact ionization of molecular hydrogen, a system that has been much less explored, both theoretically and experimentally. However, to observe the minimum in this case, it is necessary to consider ionization from molecules at specific orientations. For example, in Fig. 2 the dashed line shows the TDCS for ionization of H_2 at molecular angles $\theta_N = 50^\circ$; $\phi_N = 0^\circ$, (where the angles are with respect to the z -axis, which is defined by the incident electron direction). A sharp minimum is clearly evident. The only experiments to date on this system studied unoriented molecules, and no deep minima were found. The TDCC calculations (solid line), when averaged over all orientations, were also found to be in excellent agreement with experiment, as shown in Fig. 2. It is hoped that this study will stimulate the measurement of electron-impact ionization cross sections from oriented molecules, which is a formidable technical challenge. Further investigations of this intriguing phenomenon are also currently underway, with attention being placed on whether vortices in the correlated electronic wave functions are also playing a role.

For more information contact James Colgan at jcolgan@lanl.gov.

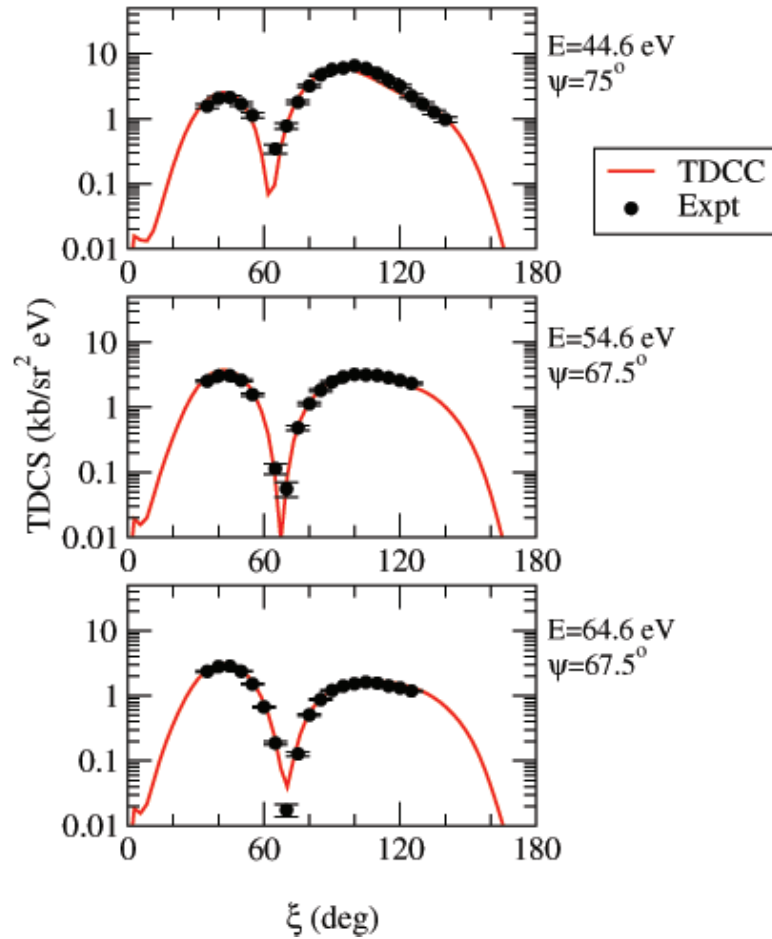


Fig. 1. Triple differential cross sections for the electron-impact ionization of helium for three incident electron energies and gun angles as indicated. In all cases the outgoing electrons have equal energy sharing. The experimental data are compared with TDCC calculations (solid red lines).

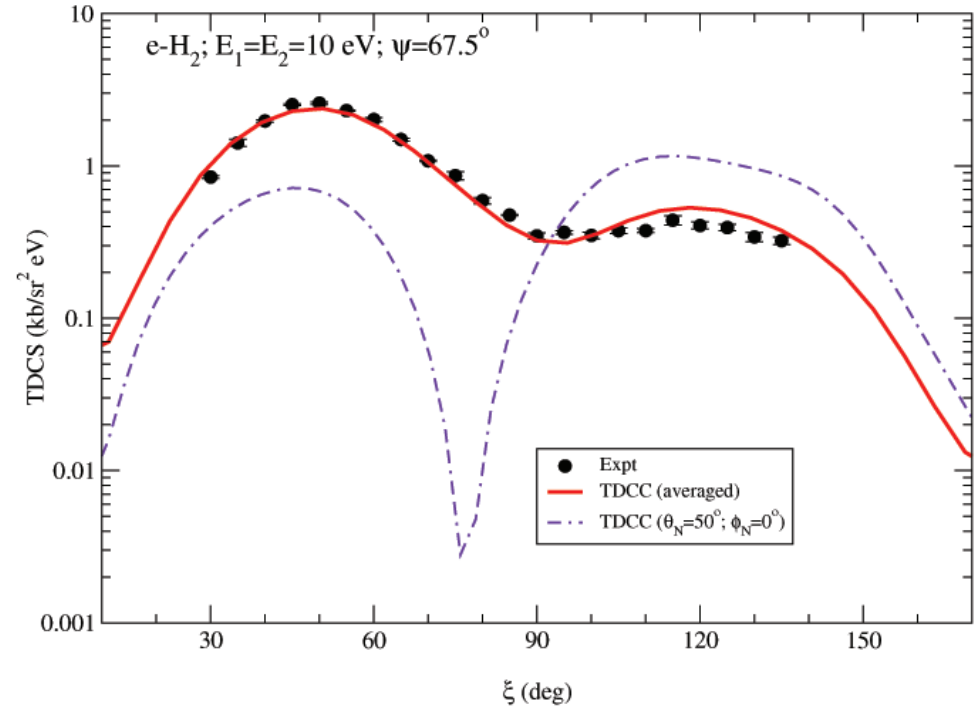


Fig. 2. Triple differential cross sections for the electron-impact ionization of molecular hydrogen at an incident electron energy of 35.4 eV. Experimental data are compared with TDCC calculations (solid red line), which are averaged over all molecular orientations. The double-dashed purple line indicates a TDCC calculation for a specific molecular orientation ($\theta_N = 50^\circ$; $\phi_N = 0^\circ$), where it can be seen that a deep minimum is predicted.

- [1] J. Colgan et al., *J. Phys. B* **42**, 171001 (2009).
- [2] A.J. Murray, F.H. Read, *J. Phys. B* **26**, L359 (1993)
- [3] A.J. Murray, F.H. Read, *Phys. Rev. A* **47**, 3724 (1993).
- [4] T.J. Reddish et al., *Phys. Rev. Letts.* **100**, 193001 (2008).

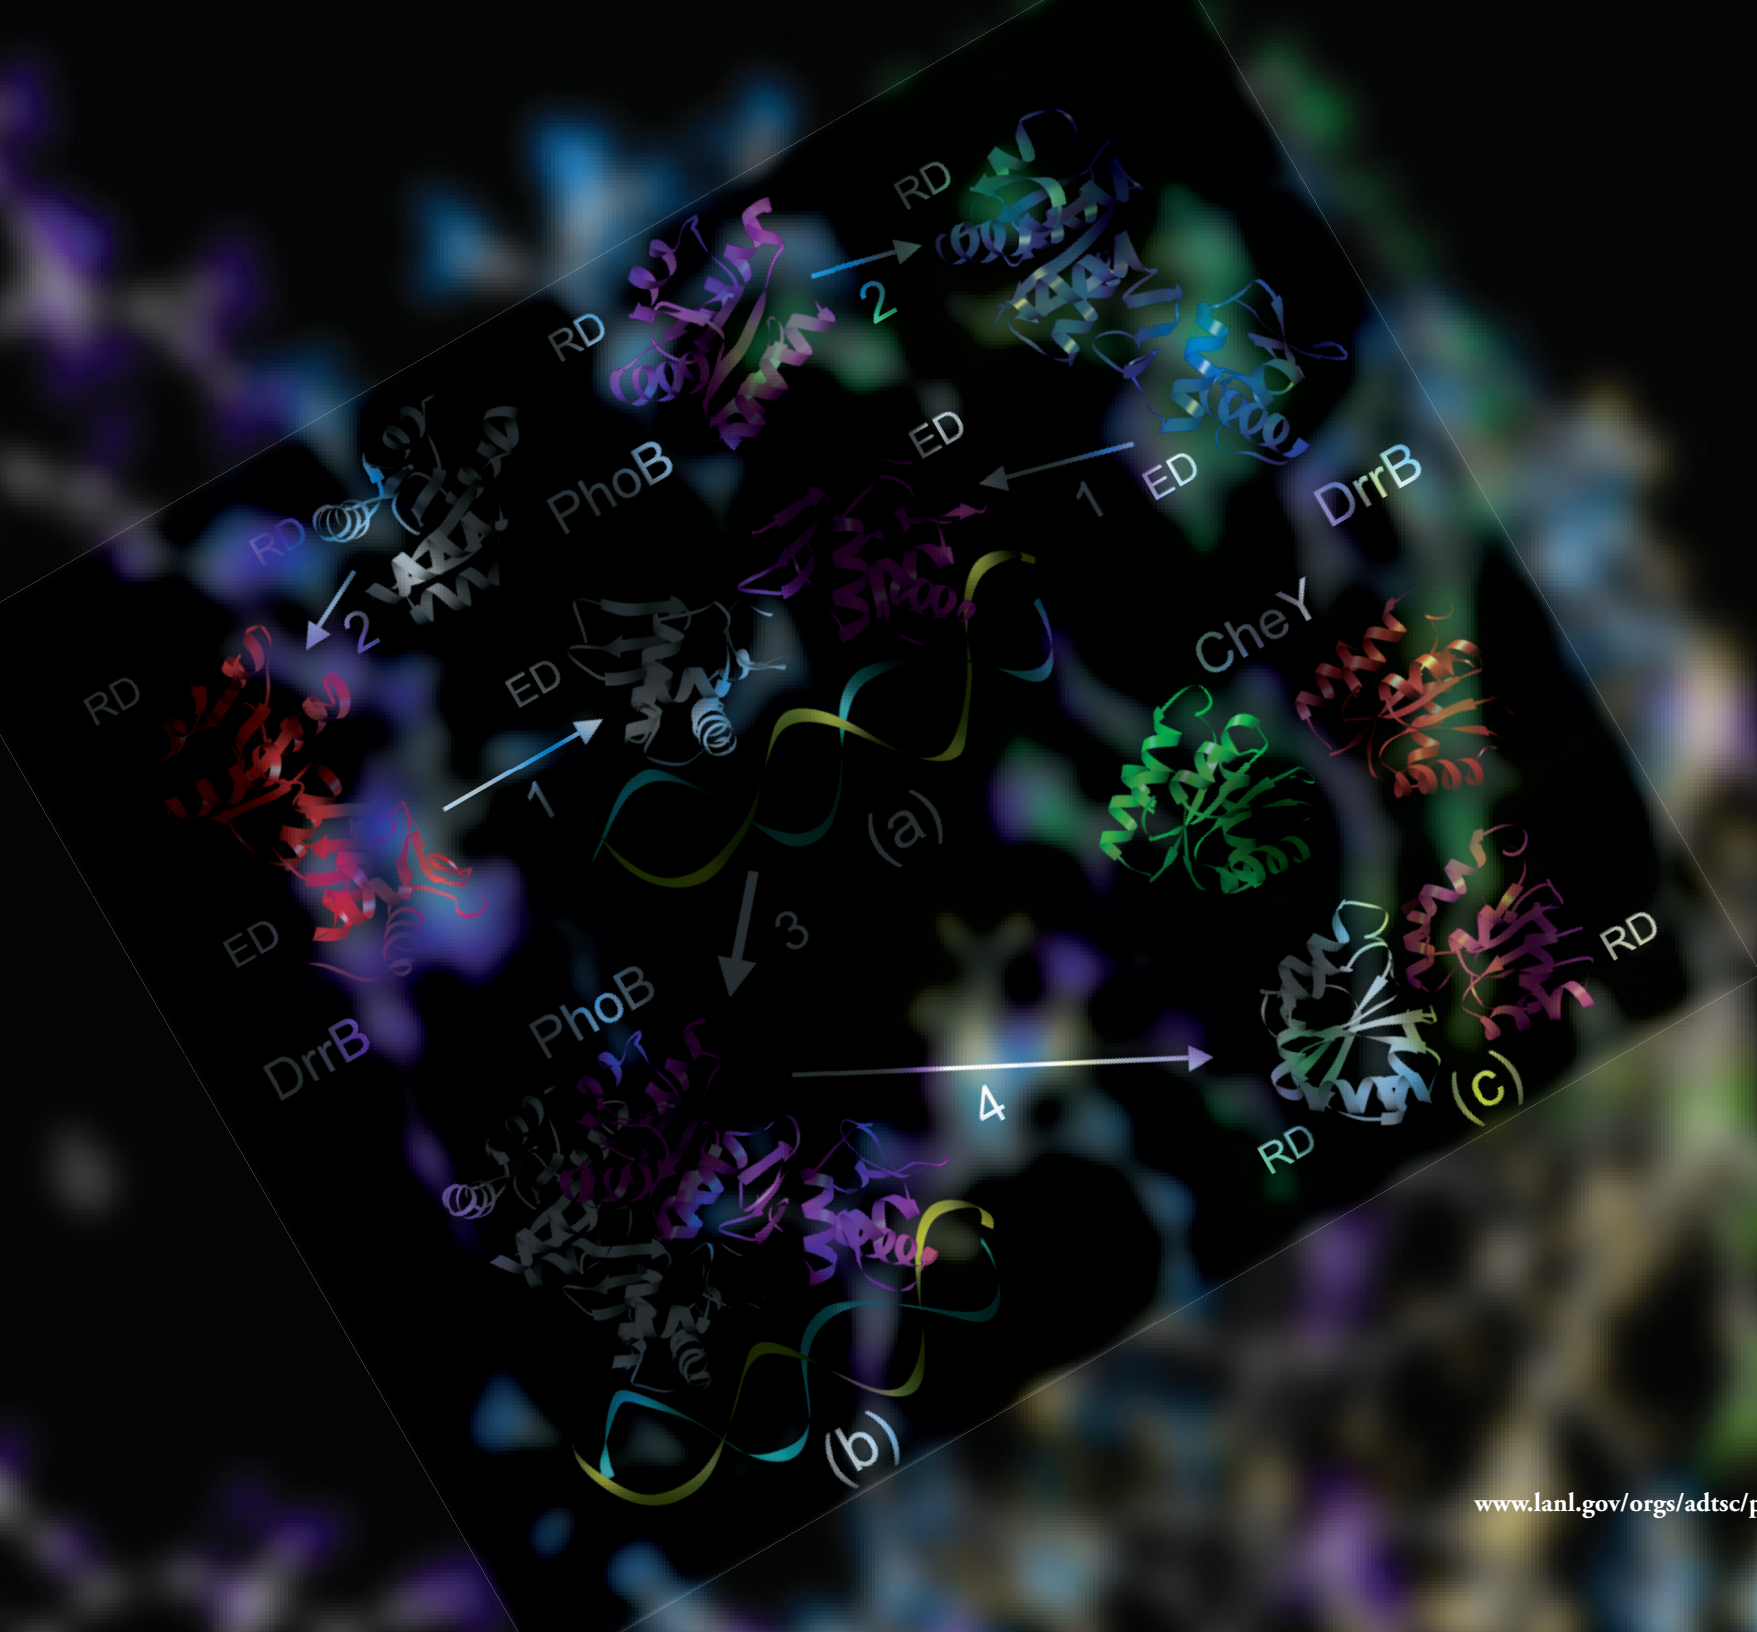
Funding Acknowledgments

- DOE, NNSA Advanced Simulation and Computing (ASC) Program
- LANL Institutional Computing Program

Chemistry and Biology

Los Alamos has a long history of outstanding work in chemistry and biology and current research in ADTSC is no exception. In this section the reader will find examples of how our scientists combine their expertise in one field to find applications in another field. For example, one of the articles explores how to use supercomputers to analyze the evolution of the HIV virus. Another explores aspects of soluble cellulose as part of a project that is studying biofuels. The functionality of the so-called “junk DNA” is also reported on. Structural biology is represented by an article on how to build models of protein complexes. The final article describes work on gene regulation using computational models.

Chemistry and Biology

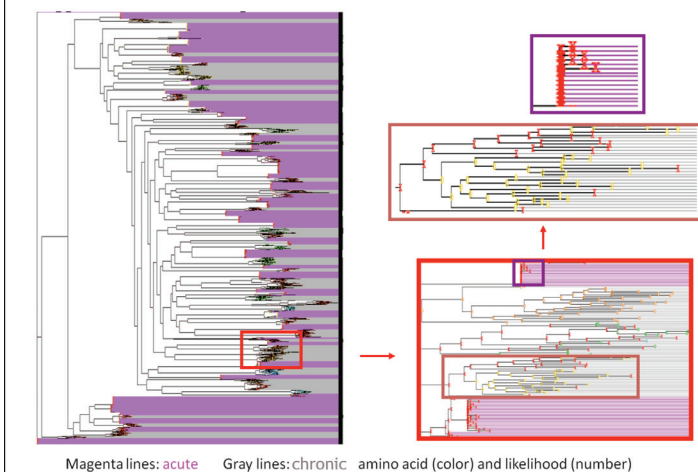


HIV vs Supercomputing

Tanmoy Bhattacharya, T-2; Marcus Daniels, S. Gnanakaran, Bette Korber, T-6

Fig. 1. Founder effects lead to entire clades sharing characteristics. If the sampling of the various clades is nonrandom, this leads to apparent correlations with traits of interest. True causal correlations, however, show up as correlations with changes.

Qualitative differences of acutes and chronics



Rapidly evolving viruses pose one of the major public health threats today. Among these, the Human Immunodeficiency Virus (HIV) that causes the Acquired Immune Deficiency Syndrome (AIDS) is particularly devastating, infecting 33 million people, with millions of AIDS-related deaths and new infections each year. Vaccines against such highly variable viruses have been unable to cope with the diversity of circulating strains: When a vaccine immunogen is presented to the human body, the elicited immune memory fails to recognize most other strains of the virus. This situation calls for both a thorough understanding of the plasticity of the virus in its war against the human host, and an intelligent design of vaccines that would provide lasting immunity against the virus. Starting with the establishment of a central sequence repository for the virus, to establishing that the virus has been circulating in humans since the early part of the 20th century [1], LANL has been at the forefront of such theoretical biology research and has contributed substantially to the field.

We also developed ideas of artificial immunogens that better capture the observed diversity of HIV strains than any natural strain can do [2,3]. While preliminary results on these artificial immunogens are sufficiently promising [4] to move them to human trials, it is desirable to advance the field of vaccine design from such data-mining techniques to biological knowledge-based approaches.

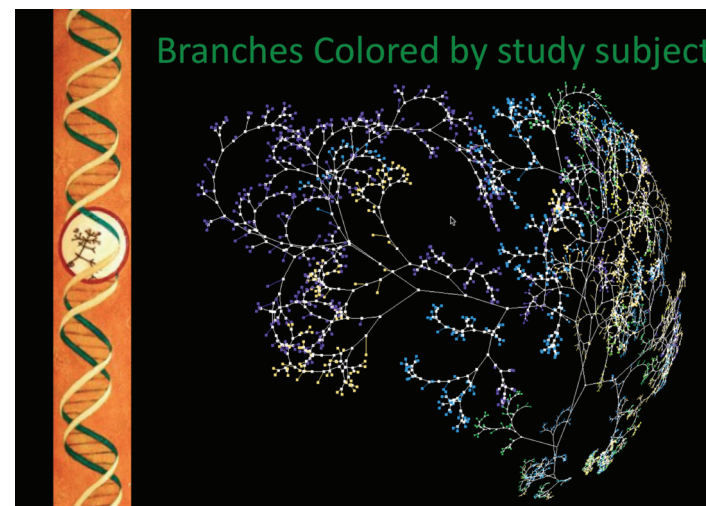


Fig. 2. A phylogeny of about 10,000 HIV sequences colored by the study subject that was used to implement phylogenetic correction on the observed correlation between genotype and phenotype.

The adaptive arm of the human immune system consists of three basic branches: 1) the Cytotoxic T Lymphocyte arm, which recognizes distinctive fragments of foreign proteins being manufactured in the body (i.e., viral proteins in infected cells) with very high specificity, 2) the Helper T Lymphocyte arm, which produces cytokines that orchestrate the immune response and have anti-viral activity, and 3) the B-cell or antibody arm, which recognizes distinctive shapes on the surface of fully folded proteins. Vast amounts of data on the interaction between HIV and all three arms of the human immune system are available, but the patterns of correlations are cryptic. Evolutionary systems are marked by long time scales, so that observed patterns in data can be due to correlations imposed by the initial historical emergence of a lineage of viruses, or *founder effects*, as well as due to biological interactions. In fact, not accounting for these effects leads to vastly erroneous statistical conclusions about the effect of the T-cell induced immunity on the evolution of the virus in populations [5]. But, whereas the sequence, i.e., the state of the virus, indeed depends on its evolutionary history, the changes that it undergoes are almost independent of changes in other lineages. Thus, true causal

correlations are also manifest in correlations with these changes (see Fig. 1), and in our work it was shown to detect effects that were validated experimentally. The separation of the two effects, i.e., a *phylogenetic correction*, thus needs access to these changes, and requires us to be able to statistically assess the genealogical relationships between the viruses and reconstruct the ancestral forms of the viruses.

Fortunately, evolution happens by the accumulation of random mutations, most of which are effectively neutral in that they do not affect the fitness of the virus to live and infect its hosts. The covariation of these mutations, then, carries a signal of shared history. This covariation can be used to construct a phylogenetic tree and an evolutionary model that leads to random changes, and, simultaneously, the ancestral forms of the virus are also reconstructed statistically. This reconstruction is technically challenging because the number of possible relationships grows factorially with the number of sequences sampled, and even heuristic searches fail to find reasonable models without extensive computations.

For example, a vaccine needs to prepare the body for fighting an incoming virus that can establish an infection in the healthy body. The virus that exists in a chronic patient, however, results from a long process of virus-host interaction and may be qualitatively different. The characterization of these differences is, however, a daunting task: the viral diversity in a chronic patient needs to be represented by at least three- to four-dozen sequences each, but to control for the phylogenetic effects, we need at least 200 to 300 patients who are infected with various subtypes of the virus. This means that one needs to fit together some 10,000 HIV sequences into a giant family tree of HIV viruses, and find those patterns that distinguish acute and chronic viruses. We therefore used Roadrunner to construct a phylogenetic tree of about 10,000 sequences from over 400 people (see Fig. 2) and are using it as a foundation to study the differences between acute and chronic sequences. This tree is currently being analyzed for pan-subtype signatures, as well as those that may be specific to individual subtypes.

We have since started analyzing data that directly measures the immunogenicity of viruses by collecting antibody-containing sera

and viruses from the same patients. These sera are seen to cluster into groups with markedly different neutralization potencies [6]. The viral envelope (Env) glycoproteins, gp120 and gp41, are the main targets of antibody (Ab) neutralization. We therefore then looked at the Env sequences of viruses from the subset of patients who make potent sera that neutralize the activity of a vast panel of viruses, and compared them against those that make average or poor responses. Even though we did not know a priori whether the difference that we observed was due to host genetics, viral factors, or stochastic events, a preliminary analysis uncovered sites in the viral sequence where changes correlated with the induction of a good immune response. Interestingly, in a 3D X-ray structure of gp120 (see Fig. 3), these sites clustered around the part of the virus that binds to the human CCR5-coreceptor, an interaction that mediates viral entry into cells, and a part of the viral envelope that had long been suspected to be involved in the induction of beneficial antibodies. In the current analysis, however, we used a small number of sera. This analysis is being extended to a much larger panel.

The advent of petaflop-scale computing, exemplified by Roadrunner, is coming to the rescue, and in the near future we expect to see a fully detailed phylogenetic analysis of such problems. Such computational techniques, complemented with our advance in experimental methods and theoretical understanding, we hope will usher in a new era that finally stops this deadly epidemic.

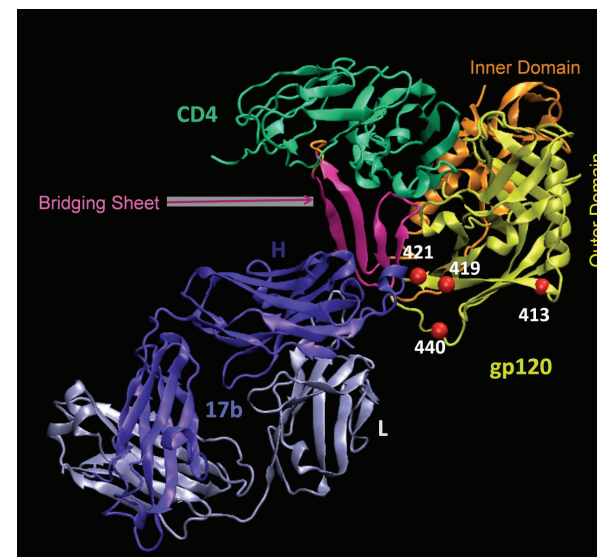


Fig. 3. Sites where mutations correlate with the induction of broadly neutralizing activity cluster in the CD4-inducible region of the HIV envelope glycoprotein gp120. When this protein (yellow, red, and orange) interacts with the cell surface molecule CD4 (green), a conformational change opens up the CD4-inducible region where our signature sites (red balls) are found. The X-ray structure (PDB code: 1RZK) also shows an antibody (blue) that binds to this region.

**For more information contact
Tanmoy Bhattacharya at
tanmoy@lanl.gov.**

- [1] B. Korber et al., *Science* **288**, 1789 (2000).
- [2] B. Gaschen et al., *Science* **296**, 2354 (2002).
- [3] W. Fisher et al., *Nature Medicine* **13**, 100 (2007).
- [4] S. Santra et al., *Proc. Natl. Acad. Sci. Unit. States Am.* **105**, 10489 (2008).
- [5] T. Bhattacharya et al., *Science* **315**, 1583 (2007).
- [6] N.A. Doria-Rose et al., *J. Virol.*, in press (2009).

Funding Acknowledgments

LANL Directed Research and Development Program

Conformational Variability of Soluble Cellulose Oligomers

S. Gnanakaran, T-6; Tongye Shen, University of Tennessee and ORNL; Paul Langan, B-8; Alfred D. French, Glenn Johnson, Southern Regional Research Center, USDA

Cellulose occurs naturally as microfibrils in the cell walls of higher plants and is also produced by some bacteria, algae, fungi, and sea animals. Two distinct crystal phases, Ia and Ib, are found in nature (collectively called cellulose I) in proportions that depend on the origin of the cellulose. For many applications, raw cellulosic material is pretreated or processed so that the cell wall architecture is disrupted and so that the cellulose crystal structure is transformed from cellulose I to other crystal phases that have improved properties for their specific application. These different crystal phases have very different properties. In particular, their different susceptibilities to digestion by cellulases is of importance in biomass conversion, i.e., the conversion of biomass into sugars that can be used in the production of biofuels. As part of our ongoing research in biomass conversion, we have been studying the molecular aspects, both structural and dynamical, of the formation of the different cellulose polymorphs and how these molecular aspects are also important for the interaction of the different polymorphs with cellulases [1].

Our most recent work concerns individual molecules of cellulose oligomers in aqueous solution [2]. Several questions remain unanswered regarding the conformational dynamics of these short chains. For example, do the glycosidic linkages have greater flexibility when freed from the restrictions imposed by the crystal environment? How do the local conformational propensities change with increasing degree of polymerization? Which internal motions are the most sensitive to thermal perturbation? It is useful to address these questions for different lengths of cellulose chains because during enzymatic depolymerization, varying lengths of cellulose chain fragments are generated. Short glucose oligomers produced by these

enzymes can hinder the overall efficiency of biomass conversion to sugar. These effects include product inhibition, competition with substrates of other enzymes, and interference with the synergetic action of cellulases by adsorption back onto the cellulose microfibril. Therefore, structural and dynamical characterization of different oligomer sizes and their relative stabilities needs to be taken into consideration when designing enzymatic cocktails.

In this study, all-atom replica exchange molecular dynamics (REMD) simulations have been used to probe the molecular aspects of the soluble oligosaccharides. The all-atom details of the simulations ensure that intra-chain and solvent hydrogen bonding, hydrophobicity and solvent ordering, and dynamics are properly taken into account. This work is one of only a few theoretical studies of the conformational dynamics of varying lengths of carbohydrate chains. Therefore, our effort will also serve as a measure of the quality of the force field and the capability of the chosen theoretical method (REMD) to provide efficient sampling of carbohydrates. To our knowledge, this work stands out as the first REMD study on polysaccharides in an aqueous environment. The use of REMD and long simulation times should provide more thorough sampling than in previous studies, and the multiple temperatures considered add considerable new information.

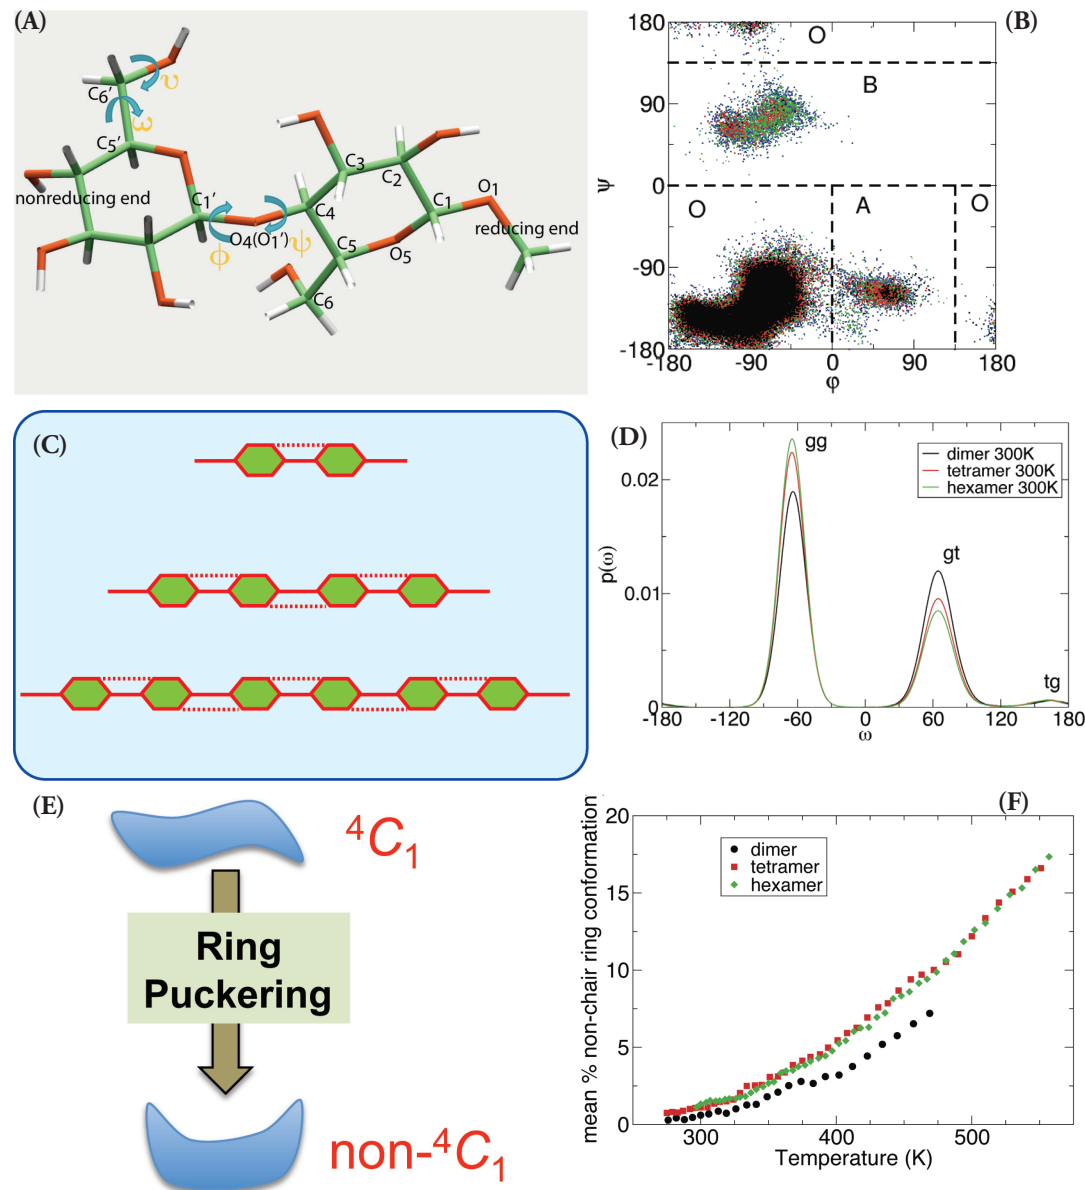
Several interesting properties are obtained from REMD simulations. As chain length is increased, the conformations of the oligomers become more rigid and likely to form intra-chain hydrogen bonds, like those found in crystals. Several other conformations and hydrogen bonding patterns distinguish these short cellulose chains from those in cellulose crystals. These studies have also addressed the key role played by solvent on shifting the conformational preferences of the oligosaccharides with respect to vacuum and crystals. Correlation between pyranose ring flipping and the conformation of the 1,4-glycosidic bond was observed. Comparisons to known measurements of oligomers and crystals are made whenever possible.

Finally, we have shown that cellulose oligomers have different properties from chains in the crystalline forms of cellulose. The results presented here may be relevant for understanding the properties of cellulose fragments during degradation by cellulases. During this enzymatic process, cellulose chains are frayed and extracted from the crystalline environments. They can potentially form alternate conformations with cellulase or when transferred from the crystalline phase to the aqueous phase.

For more information contact S. Gnanakaran at gnana@lanl.gov.

- [1] T. Shen, S. Gnanakaran, *Biophys. J.* **96**, 3032 (2009).
 [2] T. Shen, et al., *J. Am. Chem. Soc.* **131**, 14786 (2009).

Fig. 1. (A) Atomistic structure of methyl beta-D-cellobioside, representing the dimer. Torsional angles of glycosidic backbone and hydroxymethyl side chain are marked. (B) The distribution of the critical torsion angle for the glycosidic linkage in dimer is shown. The division of backbone torsion conformational space into three discrete regions O, A, and B is also displayed. (C) Schematic structure of dimer, tetramer, and hexamer considered in the simulation study. (D) The populations of the torsion angles of the hydroxymethyl side chain are plotted in terms of two-letter code that is used to describe the three conformations (gg, gt, and tg). (E) A schematic picture illustrating the possible sugar ring puckering in cellulose. Sugar rings are in chain-like (4C_1) conformation in cellulose crystals. (F) This plot captures the sugar ring puckering in cellulose oligomers. The mean percentage of glucopyranose rings in strongly distorted non- 4C_1 conformations is shown as a function of temperature.



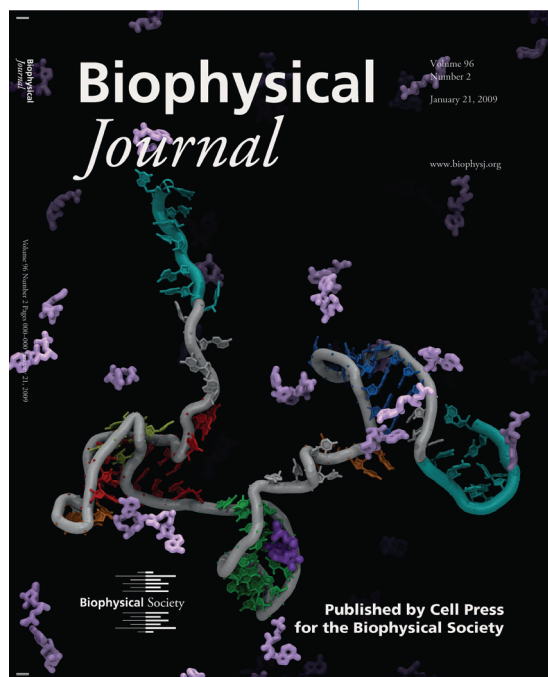
Funding Acknowledgments

LANL Directed Research and Development Program

Probing the Dark Matter of the Genome: Mechanistic Studies of Noncoding RNAs

Karissa Y. Sanbonmatsu, Scott P. Hennelly, T-6

Fig. 1. All-atom structure-based simulations performed by Paul Whitford study the interplay between RNA folding and metabolite binding to riboswitches. Gray: S-adenosylmethioine (SAM) riboswitch aptamer domain. Purple: SAM molecules surrounding riboswitch. The simulations revealed that helix P1 (cyan) forms after tertiary core collapse.



It has recently been shown that approximately 98% of the human genome is composed of DNA (deoxyribonucleic acid) that is transcribed into RNA (ribonucleic acid) but does not code for proteins. Many noncoding RNAs have been found to regulate the expression of genes. However the functions of the vast majority of noncoding RNAs have yet to be identified. One noncoding RNA that we can investigate at the molecular level is called the riboswitch. Riboswitches sense the presence of small molecules. If the small molecule is present, a gene is turned off. If the small molecule is not present, a gene is turned on. These molecules were first predicted and synthesized in the lab and later discovered in nature. Approximately 2% of

all bacterial genes are regulated by riboswitches. Furthermore, naturally occurring riboswitches have been re-engineered to detect other molecules and to control different aspects of the cell. While the 3D structures of portions of riboswitches have been determined, the mechanism of riboswitch operation is not understood. By understanding this mechanism, we will be able to customize riboswitches to operate in a desired capacity. One such application is a biosensor.

Our overall strategy is to integrate experimentation with simulations at the design level: simulations are used to design experiments and experiments are used to design simulations. For example, all-atom molecular simulations may be used to find positions on the riboswitch more suitable for fluorescent probes for experimentally monitoring conformation. Experiments reveal particular riboswitch positions that are impor-

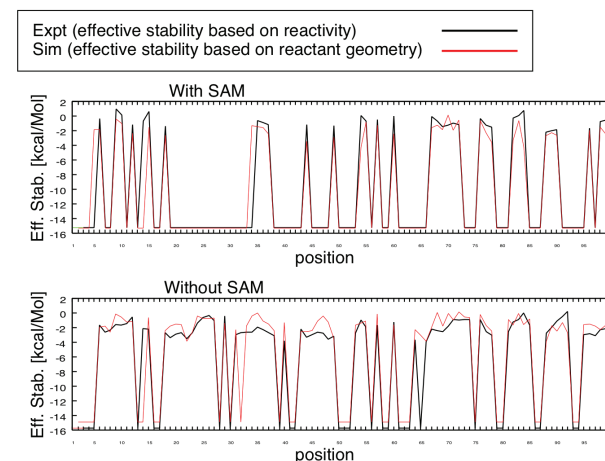


Fig. 2. Comparison between Hennelly's experiments and Hennelly's simulations for equilibrium simulations of SAM riboswitch aptamer. Black: experiment, Red: simulation. Inline probing experiments measure reactivity of RNA backbone cleavage reaction, which indicates relative mobility of backbone. A fitness function based on the geometry of the cleavage reaction is computed from simulations.

tant for metabolite binding. These positions are then studied in detail in the simulations to elucidate a mechanism. In another example, experiments producing melting curves are performed that can be replicated in a simulation. The experimentally determined melting curves then dictate which temperatures to simulate and which force field parameters to use.

As a first step into understanding riboswitch operation, Paul Whitford (T-6) and collaborators at University of California at San Diego performed the first all-atom simulation of a riboswitch aptamer domain (Fig. 1). The aptamer domain is the portion of the riboswitch that binds the small molecule. The expression platform is the portion that turns the gene on or off. This aptamer domain has four helices, P1, P2, P3, and P4. The conventional wisdom in RNA folding states that secondary structure (base pairing and individual helix formation) should form before tertiary structure (contacts between helices). While this sequence of events holds true for most of the helices during our simulation, we find that a functionally important helix, P1, forms after much of the tertiary structure. This occurs because helix



Fig. 3. Caliper liquid-handling robot purchased by Scott Hennelly enables high-throughput inline probing experiments. Reactions are performed on the blue deck. The head unit and gripper (blue and gray block above deck towards right) moves plates around. Each plate allows 96 simultaneous reactions. Twister II arm (right-hand side) moves plates from storage to deck.

P1 requires the tertiary contacts to bring the complementary bases together. This has large implications for how the regulatory decision, on or off, is made.

Our next step was to achieve agreement between simulation and experiment to give us baseline simulation parameters that we can use in future mechanistic studies. Here, Scott Hennelly performed biochemical probing experiments that reveal the secondary structure of the RNA, and also information regarding dynamics of the RNA. The experiments (inline probing and selective 2'-hydroxyl acylation and primer extension (SHAPE probing) monitor reactions along the backbone of the RNA that reflect the mobility of the backbone at each base position along the RNA. For example, nucleotides in helices are shown to be less mobile, while nucleotides in loops and in junctions between helices are more mobile. Hennelly then performed simulations and calculated mobilities of each backbone position. Next a script was written to automatically adjust force-field parameters to achieve better agreement between simulation and experiment. This achieved close agreement between simulation and experiment with a few outlier nucleotides. The parameters for these outlier nucleotides were then individually tuned, resulting in Fig. 2. To understand the entire riboswitch, including the aptamer domain

and the expression platform, Hennelly developed a new experimental assay that monitors switching between the two riboswitch states. By using a two-piece system of RNA instead of only one piece, he is able to observe the switch transitioning from the off state (terminator formed) to the on state (antiterminator formed). Specifically, a fluorescent label molecule is placed on the second strand. If it combines to form the antiterminator, the fluorescence is quenched. Our main result is the discovery that 3D interactions in the sensor domain of the riboswitch are essential for proper riboswitch operation. During normal operation, there is a large difference between switching rates with and without the small molecule. When Hennelly removed 3D interactions by mutating the riboswitch sensor domain, he found that there is no longer a difference. Therefore, the 3D interactions are very important and dependent upon the binding of the small molecule. Hennelly and Sanbonmatsu are both Principal Investigators on a recently funded NIH grant to apply this method to riboswitches.

By performing the first study of the switching mechanism during riboswitch operation, we have laid the groundwork for designing customized biosensors with user-specified switching rates. Often a tradeoff occurs between the sensor domain and the switching domain of the biosensor. We find that 3D interactions dramatically affect this tradeoff. This effect was not known before and will have a large impact on noncoding RNA biosensor design. Designing effective biosensors is the first step to next-generation bioremediation (allowing us to track down hazardous waste) and to more effective chemical agent detectors. The current work not only paves the way for the design of biosensors with novel affinities, but also sensors with low false positives and increased signal. There are also applications in bioenergy, where riboswitches can be used to detect products of cellulose degradation, allowing the engineering of more efficient cellulases (enzymes that break down cellulose).

In the future we plan to use high-throughput platforms and next-generation sequencing methods to design new and more complicated types of RNA switches. Hennelly purchased an automated liquid handling robot (Fig. 3) from LDRD capital equipment funds and is

using it for high-throughput secondary structure and RNA dynamics determinations. We plan to use this system, in combination with Illumina sequencing, to perform a large-scale molecular evolution study of RNA switches. In standard molecular evolution, or in vitro selection, RNAs are randomly mutated and more efficient mutants are selected. Successive rounds are performed to optimize the desired characteristics of the RNA. We will use Illumina technology to determine the entire fitness landscape in a single round with much more sampling of sequence space. The promise of riboswitches as either sensors or regulators is in their high specificity and ability to adopt two dramatically different conformations based on binding a small molecule. We have developed a ground-breaking methodology to simultaneously select for these two important characteristics.

For more information contact
Karissa Y. Sanbonmatsu at
kys@lanl.gov.

Funding Acknowledgments

- LANL Directed Research and Development Program
- NIH, American Recovery & Reinvestment Act (ARRA)

Building Models of Protein Complexes using the Known Protein-protein Interfaces: Structural Model of PhoB Dimeric Complex in its Active State

Chang-Shung Tung, T-6

Often, multimolecular complexes are the functional forms of proteins inside cells. To understand how proteins function, it is necessary to know how they interact with each other. Knowing the structures of protein complexes is critical. It is well known that solving the structure of complexes is significantly more difficult than solving the structure of individual proteins. While homology modeling is a promising approach in modeling protein structures, it has not been proved successful in modeling protein complex structures. Improving the knowledge of protein complex structures remains a great challenge. The good news is that a wealth of structural information exists. In the current version of Protein Data Bank (PDB) it can be seen that many of these proteins exist in a multimer conformation that contains the information of interaction surfaces (interfaces). Here we present a novel approach that utilizes the information of known protein interfaces to build models of protein complexes. This approach will significantly improve protein complex structural modeling and prediction.

The protein complex that is used to illustrate the approach is the PhoB dimer in its functional active state, which binds to the targeted DNA duplex. In the crystal structure (PDB accession code: 1GXP), two PhoB EDs (magenta and white molecules) bind to the targeted DNA direct repeat pho box (cyan and yellow molecules) as shown in Fig. 1a. The binding of DNA direct repeats makes the ED dimer follow a tandem symmetry. However, the known structure of the PhoB RD dimeric complex (PDB accession code: 2JB9) follows a rotational symmetry. It is not obvious how to best dock a RD dimeric complex with a rotational symmetry onto an ED dimeric complex

with a tandem symmetry. A number of response regulators exist with their two-domain structures solved experimentally (PDB accession codes: 1KGS, 1P2F, 1YS6, 2GWR, 2OQR, 1A04, 1YIO). These structures contain the information of ED/RD interfaces. Applying the information together with the structure of the ED/ED dimeric complex (1GXP), we explore the potential solutions for the PhoB dimeric complex. Out of the ED/RD conformations, only that of DrrB (1P2F), shown as the red and the blue molecules in Fig. 1), a PhoB/OmpR homolog, provides a satisfactory solution where the two RDs are in contact but not overlapping. Combining the structural information of ED/ED (1GXP), ED/RD (1P2F), ED (1GXP) and RD (2JB9), the model of the PhoB dimeric complex is developed (shown as the white/magenta molecules bound to DNA in Fig. 1b). This model structure has appealing features including: good stereochemistry (no clashes between domains, stable interface surface area), protein-like structure (contents of secondary structures, density, etc.), and several of the known interaction interfaces. However, the RD/RD dimer in the modeled complex follows a tandem symmetry that is different from the rotational symmetry observed in the PhoB RD/RD dimer (2JB9). The remaining question is: Does the new interface between the two RDs in the model exist in other protein domains of a similar fold? To answer this question, we search for interfaces between domains that have the flavodoxin-like fold and give the two domains a tandem symmetry. Interestingly, the dimeric structure (PDB accession code: 1FFY) of CheY, a chemotaxis protein, has two flavodoxin-like domains following a tandem symmetry. This dimeric structure of CheY (1FFY) is very similar to that of the PhoB dimeric RDs as shown in Fig. 1c. We have demonstrated that the information of protein interfaces and folds can be directly used to build structural models of protein complexes.

For more information contact Chang-Shung Tung at ct@lanl.gov.

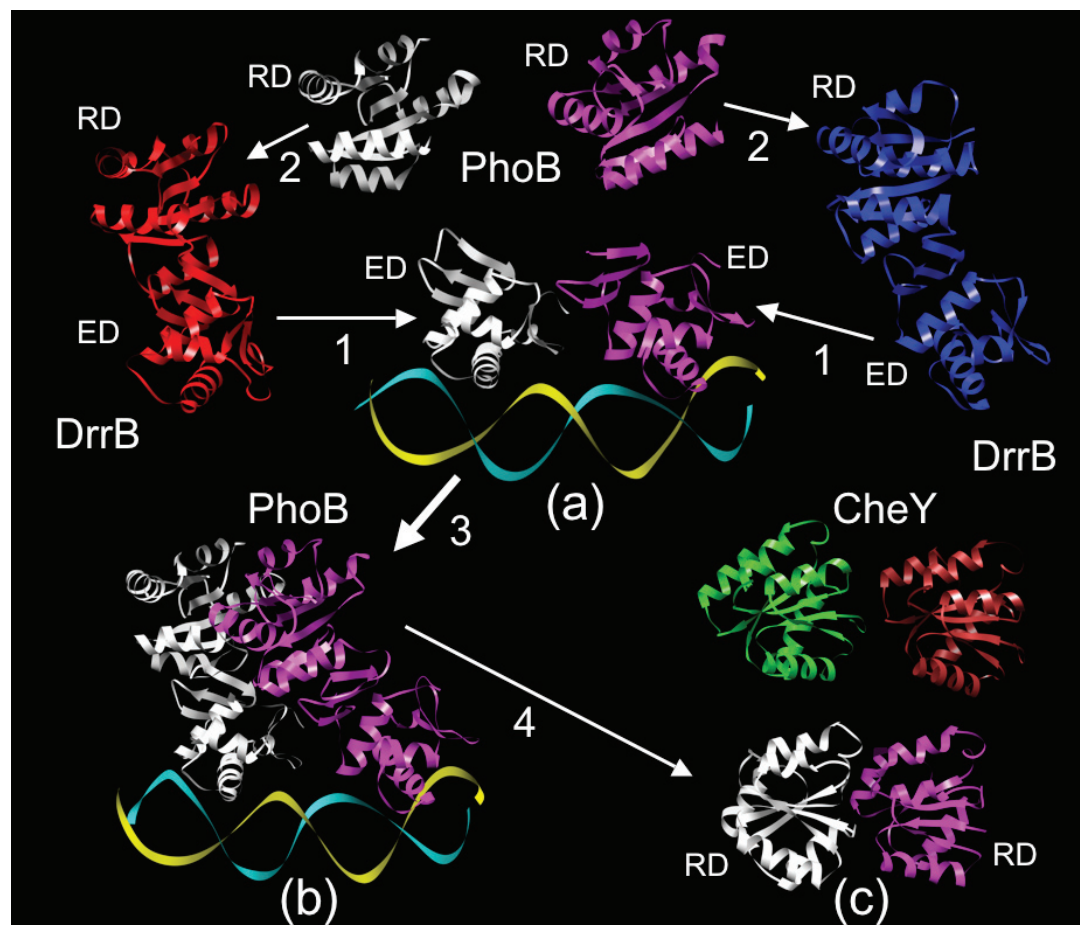


Fig. 1. Structural model of the PhoB dimeric complex as it binds to its targeted DNA duplex. The matching ED of DrrB to ED of PhoB is shown in 1a. The resultant PhoB/DNA model that is shown in 1b. The RD/RD interface in CheY is similar to that in the model as shown in 1c.

Funding Acknowledgments

LANL Directed Research and Development Program

Quantitative Insights into Gene Regulation

Robert G. Martin, Judah L. Rosner, National Institutes of Health;
Michael E. Wall, CCS-3

When environmental conditions change, cell survival can depend on the sudden production of proteins that are normally in low demand. Protein production is controlled by transcription factors that bind to DNA near genes and either increase or decrease RNA production. Many puzzles remain concerning the ways transcription factors do this.

There are hundreds of transcription factors in *Escherichia coli* and while most of these target only a small number of genes, there are several that regulate expression of 10 or more genes. Taken together, such global transcription factors directly regulate more than half of the ~4,300 genes in *E. coli*, and their regulatory interactions yield important insights into the organization of the genetic regulatory network [1-3]. Because they regulate so many genes, global transcription factors also play a large role in controlling cellular behavior; however, insights into behavior are currently limited by a lack of quantitative information about how transcription factors differentially regulate target genes.

We are performing quantitative studies of gene regulation by the global transcription factor MarA in *E. coli*. MarA activates ~40 genes of the *E. coli* chromosome, resulting in different levels of resistance to a wide array of antibiotics and superoxides (see [4] for references). We placed the expression of MarA under the control of an external signal and examined the expression of 10 target genes as a function of activator concentration [5]. We found that activation of different genes occurs in a well-defined order with respect to the level of MarA, enabling cells to mount a response that is commensurate to the level of threat detected in the environment. In contrast with a commonly

held assumption, we found that the order of activation does not parallel the strength of MarA binding to DNA near the target gene. This finding suggested that interactions between MarA and the RNA polymerase transcriptional machinery play an important role in determining the order of activation, but the data did not immediately reveal what the nature of these interactions might be.

Next, we developed a computational model of gene regulation to understand how interactions between MarA and polymerase activate transcription of *marRAB*, *sodA*, and *micF* [6]—of the 10 targets we examined previously, these three are the only ones that exhibited saturation at high MarA, which provided an important constraint for the modeling. The model was specifically designed to compare the textbook model of recruitment, in which MarA increases polymerase binding but does not increase the rate of post-binding events [7,8], with a more general model in which activator can either increase or decrease polymerase binding, and can either increase or decrease the rate of post-binding events (Fig. 1). The model clearly explains the lack of correspondence between the order of activation and the MarA-DNA affinity, and indicates that the order of activation can only be predicted using information about the strength of the full MarA-polymerase-DNA interaction. Instead of favoring activation by recruitment, the modeling favors activation by increasing the likelihood that polymerase will initiate transcription when it is already bound to DNA near the target gene (Fig. 2). It also suggests that MarA can activate expression while decreasing the overall presence of the transcription machinery at the start of a gene. This mechanism is opposite to the textbook model of recruitment; nevertheless it enables cells to respond quickly to environmental challenges and is likely of general importance for gene regulation in *E. coli* and beyond.

For more information contact Michael E. Wall at mewall@lanl.gov.

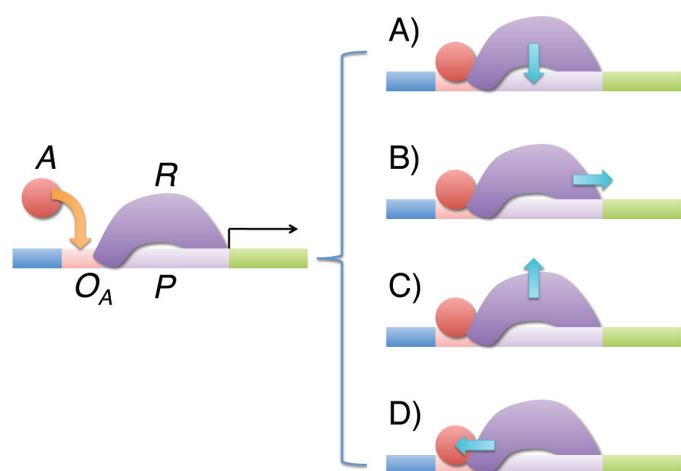


Fig. 1. Illustration of ways that bacteria might control mRNA transcription. A) In the textbook model of activation by recruitment, an activator (A) binds to an operator sequence (O_A) and increases the presence of RNA polymerase (R) at the start site (P). B) Although usually ignored, it has long been known that activators can accelerate the initiation step when polymerase is bound. C) Wall et al. [6] found that activation can involve a decrease in the presence of polymerase at the start site; although this effect is naturally associated with B), it was previously ignored. D) Activator can also retard the initiation step, a neglected effect that is naturally associated with the textbook model.

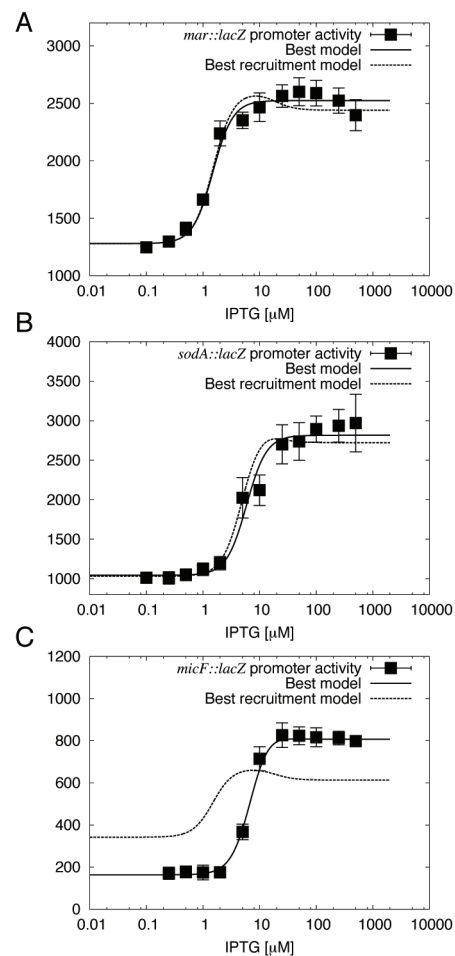


Fig. 2. Comparison of alternative models of transcriptional activation. A) For activation of *marRAB* the fit of the recruitment model is reasonable (Fig. 1A), but the fit of models involving acceleration of the initiation step is better (Fig. 1B). B) For activation of *sodA*, the best models involve not only acceleration (Fig. 1B), but also a decrease in the presence of polymerase at the start site (Fig. 1C). C) Activation of *micF* is poorly fit by the recruitment model, but is well fit by models that involve both acceleration of the initiation step (Fig. 1B) and a decrease in polymerase binding (Fig. 1C).

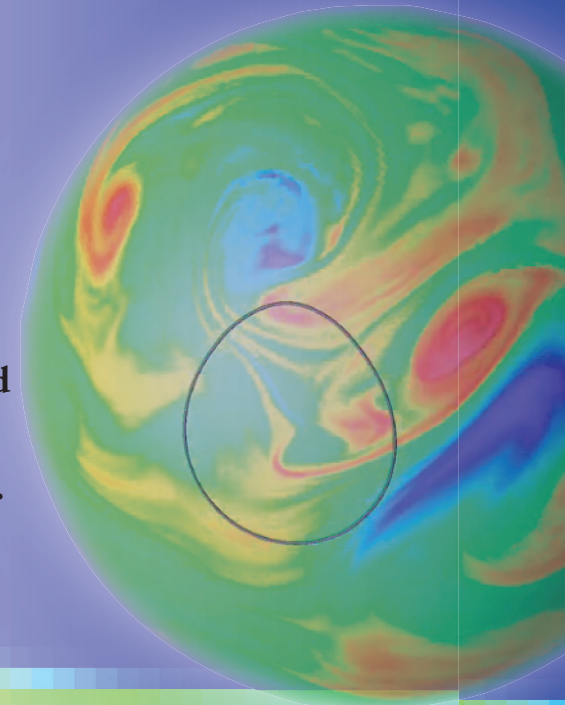
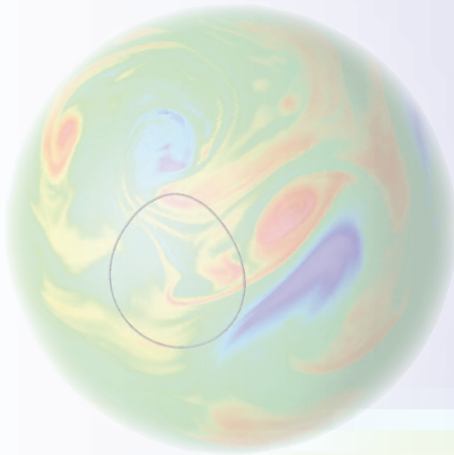
- [1] J.A. Freyre-Gonzalez et al., *Genome Biol.* **9**, R154 (2008).
- [2] A. Martinez-Antonio, J. Collado-Vides, *Curr. Opin. Microbiol.* **6**, 482 (2003).
- [3] S.S. Shen-Orr et al., *Nat. Genet.* **31**, 64 (2002).
- [4] D.G. White, M.N. Alekshun, P.F. McDermott, *Frontiers in Antimicrobial Resistance*, ASM Press (2005).
- [5] R. G. Martin et al., *J. Mol. Biol.* **380**, 278 (2008).
- [6] M.E. Wall et al., *PLoS Computational Biology* **5**, e1000614 (2009).
- [7] M. Ptashne, *Philos. Transact. A Math. Phys. Eng. Sci.* **361**, 1223 (2003).
- [8] M. Ptashne, A. Gann, *Genes & Signals*, Cold Spring Harbor Press (2002).

Funding Acknowledgments

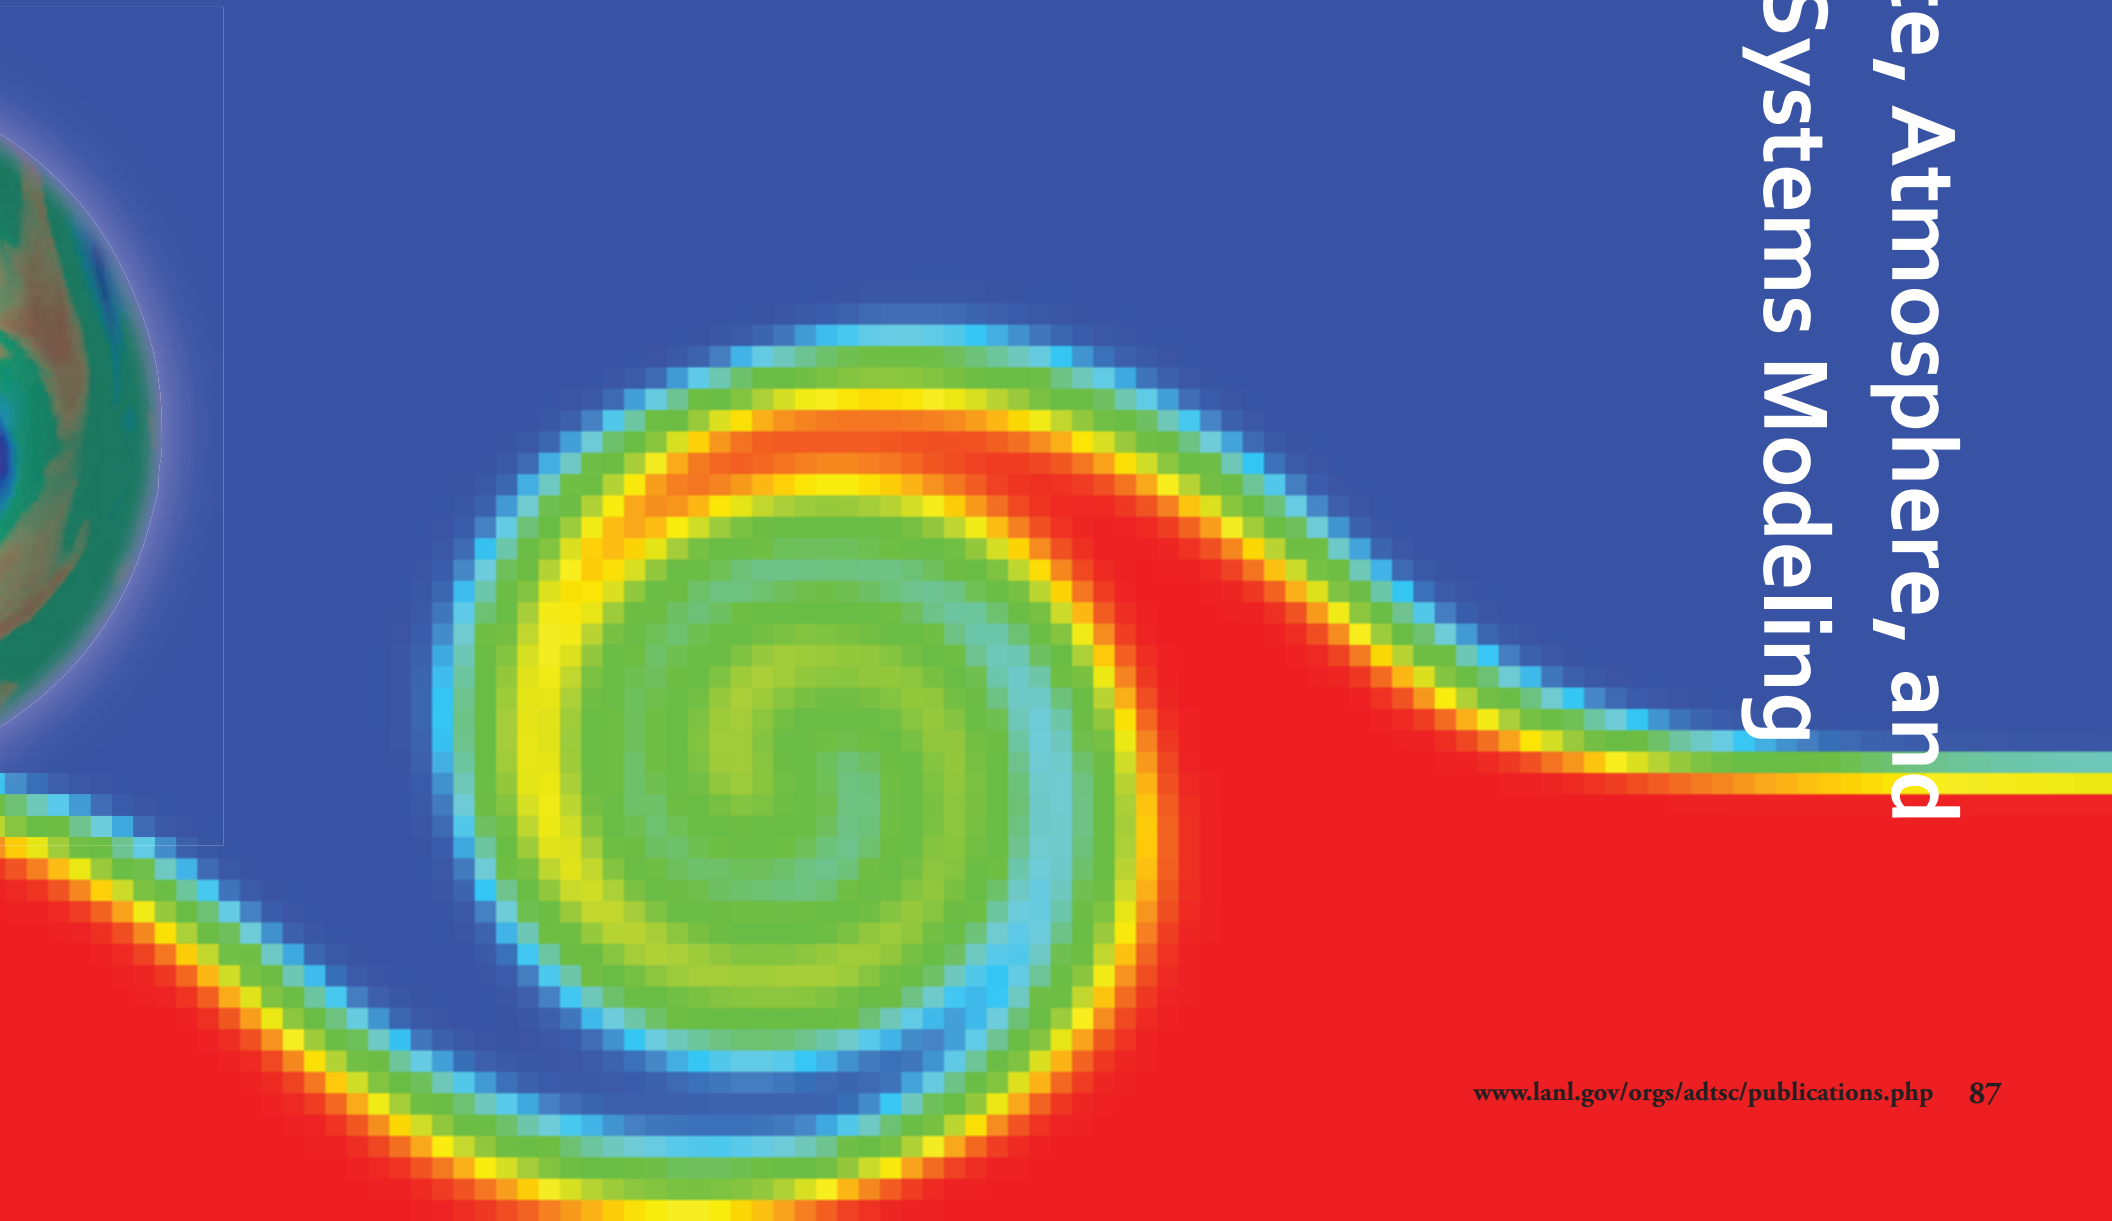
- LANL Directed Research and Development Program
- National Institutes of Health (NIH)–Intramural Research Program

Climate, Atmosphere, and Earth Systems Modeling

LANL's recent work in modeling the global and regional climate, the atmosphere and earth systems has drawn on our considerable expertise in computational science, both algorithms and leading-edge computing resources. The breadth of this work is captured in diversity of the following articles. The first article describes the observation and prediction of changes in Greenland's ice surface mass. The second article discusses the application of advanced numerical methods to study the dynamics of tracers in oceanic and atmospheric flows. Next is a description of parallel simulation of hydrological systems in semiarid regions and the final article ties together disparate length and timescales by multiresolution modeling.



Climate, Atmosphere, and Earth Systems Modeling



Observing and Predicting Changes in Greenland's Surface Mass Balance

Sebastian H. Mernild, CCS-2; William H. Lipscomb, T-3

Box 1: SnowModel

The GrIS surface water balance components, with emphasis on precipitation, evaporation, sublimation, surface mass balance, and runoff, were simulated using SnowModel, a spatially and temporally distributed meteorological and snowpack modeling system, for present-day and future conditions [6,11]. SnowModel is composed of five submodels:

- 1) MicroMet defines the meteorological forcing conditions,
- 2) EnBal calculates the surface energy exchanges,
- 3) SnowPack simulates mass and heat transfer processes due to, e.g., retention and internal refreezing,
- 4) SnowTran-3D is a blowing-snow model that accounts for snow redistribution, and
- 5) SnowAssim is a snow-data assimilation model that adjusts the simulated snow distributions towards available snow observations.

The Greenland Ice Sheet (GrIS) is the Northern Hemisphere's largest terrestrial area with a perennial ice and snow cover. Complete loss of the ice sheet would raise global sea level by about 7m [1]. Climate models suggest that a regional temperature increase of 3°C or more, which is predicted to occur during this century under many greenhouse gas emissions scenarios, would be sufficient to melt most of the GrIS during a period of centuries to a few millennia [2]. Simulations by atmosphere-ocean models for areas north of 60°N project an increased mean surface air temperature of 2.5°C by the mid 21st century and 4.5°–5.0°C by the end of the century [3].

The GrIS gains mass from surface snow accumulation and loses mass by two main processes: (1) meltwater runoff (mostly as a result of summertime surface ablation), and (2) iceberg calving (associated with the flow of large outlet glaciers into the ocean). Recent observations indicate that the GrIS responds more quickly to climate perturbations than previously thought, especially near the margin in southern Greenland [4]. The ice sheet was close to equilibrium during the relatively cool 1970s and 1980s, but lost mass rapidly as the climate warmed in the 1990s and 2000s, with no sign of deceleration [5]. The GrIS is currently losing mass at a rate of nearly 300 Gt/yr [6, 7], enough to raise global sea level by about 0.8 mm/yr. Both surface melting and glacier outflow have accelerated during the past decade, likely in response to atmospheric and oceanic warming. Ice has thinned along the periphery (primarily in the south), while ice in the interior has thickened slightly because of increased snowfall [8].

In order to predict the impact of climate changes on the GrIS, it is essential to establish the present-day surface mass balance (SMB, the difference between annual accumulation and ablation) and

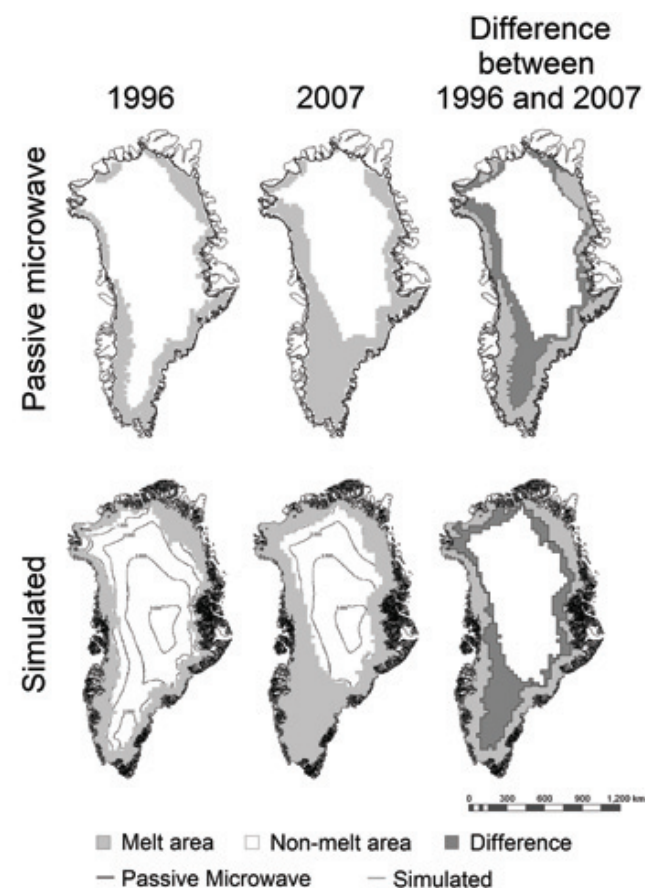


Fig. 1. Surface melt extent of the Greenland Ice Sheet based on passive microwave satellite-derived observations and SnowModel simulations for the years 1996 and 2007. The plots in the third column show the difference between the two years. The satellite-derived melt extent is based on data provided by the Cooperative Institute for Research in Environmental Sciences (CIRES), University of Colorado at Boulder.

surface melt extent (the area that experiences some melting in a given year). Relatively modest temperature changes can bring about large changes in melt volume and extent. Since 1995, the year with the minimum GrIS surface melt extent was 1996, while the most extensive melting occurred in 2007 (Fig. 1). The 2007 melt extent was 20% greater than the average for 1995–2006. In regional

modeling simulations [9], the year-to-year variability in melt extent was in excellent agreement with observations. The year 2007 had the highest simulated surface runoff (approximately 520 Gt/yr) and the lowest surface mass balance (close to zero). Through the simulation period, the SMB varied from -5 (2007) to 310 Gt/yr (1996), averaging 124 (± 83) Gt/yr (Fig. 2).

Scientists in the Climate, Ocean and Sea Ice Modeling (COSIM) project are using global and regional models to predict future changes in Greenland's surface mass balance. Mernild and collaborators [10] have used projections from a regional atmosphere model to force SnowModel, a state-of-the-art meteorological and snowpack modeling system. (See **Box 1** for further information). The GrIS surface mass balance was simulated to decrease by ~ 300 Gt/yr between 1950 and 2080, with negative values by the final decade of the simulation. (When the SMB is negative, the ice sheet melts inexorably, even in the absence of iceberg calving.) The end-of-summer melt extent nearly doubled during this period, with the greatest changes in the southern and eastern parts of the ice sheet.

These simulations are now being extended to the Community Climate System Model (CCSM), a fully coupled global climate model. COSIM scientists have recently coupled the Glimmer ice sheet model to CCSM in order to simulate changes in ice sheet dynamics. Glimmer is forced with the surface mass balance computed by the CCSM land and atmospheric models. The global atmosphere model is run at relatively coarse grid resolution (~ 100 km), which may not be sufficient to obtain an accurate SMB in regions of steep topography. To improve the estimated mass balance, we are taking part in coupled global climate-change simulations with an atmosphere grid resolution of ~ 25 km over Greenland. We are also coupling surface processes to the dynamic ice-sheet model in order to simulate melt-induced increases in ice flow. The coupled models will be validated against recent observations and then applied to various greenhouse forcing scenarios on decade-to-century time scales. These simulations will constitute a unique contribution to ice-sheet and sea-level

projections in the next Intergovernmental Panel on Climate Change (IPCC) assessment report, AR5.

For more information contact

Sebastian H. Mernild at mernild@lanl.gov.

- [1] P. Lemke et al., in *Climate Change 2007: The Physical Science Basis*, (Cambridge University Press) (2007).
- [2] J.M. Gregory et al., *Nature* **428**, 616 (2004).
- [3] R.B. Alley et al., in *Climate Change 2007: The Physical Science Basis*, (Cambridge University Press) (2007).
- [4] I. Velicogna, J. Wahr, *Geophys. Res. Lett.* **32**, L18505 (2006).
- [5] E. Rignot et al., *Geophys. Res. Lett.* **35**, L20502 (2008).
- [6] I. Velicogna, *Geophys. Res. Lett.* **36**, L19503 (2009).
- [7] S.H. Mernild et al., *J. Hydrometeorology* **9**, 1191 (2008).
- [8] W. Krabill et al., *Geophys. Res. Lett.* **31**, L24402 (2004).
- [9] S.H. Mernild et al., *Hydrolog. Process.* DOI: 10.1002/hyp.7354 (2009).
- [10] S.H. Mernild et al., *J. Hydrometeorology*, **11**, 3 (2010).
- [11] G.E. Liston, K. Elder, *J. Hydrometeorology* **7**, 1259 (2006).

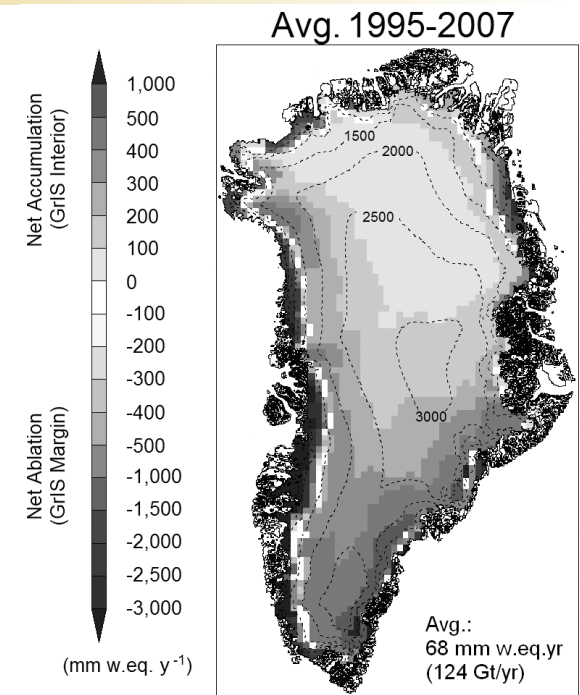


Fig. 2. The 1995–2007 average annual surface mass balance as modeled for the Greenland Ice Sheet (in mm w.e./yr and Gt/yr). Dotted lines show surface elevation with a 500-m contour interval.

Funding Acknowledgments

- LANL Institute for Geophysics and Planetary Physics
- DOE Office of Biological & Environmental Research (OBER), Climate Change Prediction Program
- DOE Scientific Discovery through Advanced Computing Program (SciDAC)

Tracer Advection for Ocean and Atmospheric Flows

Robert B. Lowrie, CCS-2; Todd D. Ringler, T-3

The physical description of ocean and atmospheric flow dynamics includes numerous tracers, which are defined as quantities that move with the flow velocity. This movement is commonly referred to as advection. Examples of tracers include flow temperature, salinity, individual chemical species, and bio-geochemical species such as plankton. These various tracers play a critical role in climate modeling, and accurately tracking their advection is an ongoing research effort.

We have recently developed a new numerical method that models tracer advection more accurately than existing methods. We refer to the method as characteristic discontinuous Galerkin (CDG).

The CDG method is based on previous work for advection using discontinuous Galerkin [1] and methods that use a geometric remap [2]. We have shown that in simple cases, the method is related to Van Leer's "exact evolution with least-squares projection" [3] and Prather's moment method [4]. The CDG method has the following five properties:

Accuracy—Many numerical methods for advection suffer from artificial mass diffusion, in that the method artificially smears and disperses the tracer field as it moves. Reducing this diffusion is critical for accurate climate modeling [5]. CDG minimizes artificial diffusion by using a high-order polynomial representation of

Fig. 1. Initial condition for the cyclogenesis problem, showing blue and red tracers.



Fig. 2. CDG results at the final simulation time. On this plotting scale, these results are indistinguishable from the exact solution. CDG used 128 mesh cells in each spatial direction and represented the solution as a cubic polynomial in each cell.

the tracer in each mesh cell. The polynomial order may be increased locally in regions where the tracer variability is large.

Generality—CDG is designed to work on general mesh topologies, such as Voronoi meshes that will be used in future ocean and atmospheric models [6].

Conservation—This means that whenever physically appropriate, the total amount of a given tracer remains constant, to within round-off of the computer floating-point precision.

Positivity—CDG enforces the known physical bounds on tracer quantities, such as that the temperature should be nonnegative. The method enforces positivity without decreasing the formal order-of-accuracy.

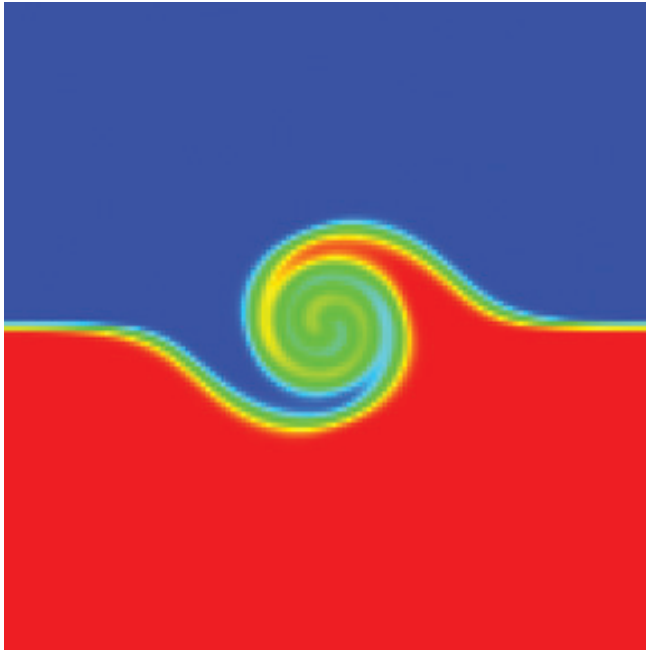


Fig. 3. Results from a legacy advection scheme; compare with Fig. 2. Unlike the results of Fig. 2, the tracer quantities here are smeared significantly across the vortex.

Speed—Tracer advection can be a dominant fraction of the computation time in certain simulations. The computational cost of CDG is low for three reasons:

- The majority of the cost is independent of the number of tracers. Geophysical flows may require tracking tens or even hundreds of tracers. CDG precomputes much of the time-step update based on the given velocity field, so that increasing the number of tracers has a negligible increase on the computation time.
- CDG parallelizes extremely well, with minimal neighboring information required for each update, and is designed to work well with advanced computer architectures such as Roadrunner. For a single layer of ghost cells, only a single parallel communication step is required per time step.

- The maximum allowable time step is independent of the polynomial order used in each cell, so that large time steps may be taken. Most other discontinuous Galerkin methods require reducing the time step as the polynomial order is increased, which results in increased computation time.

Our plan is to incorporate CDG into both ocean [7] and atmospheric flow [8] simulation computer codes. Figures 1-3 show results for the cyclogenesis test problem [9]. The initial condition for this problem consists of two tracers, separated by a line, as shown in Fig. 1. A vortex is centered in the domain, which causes the interface between the two tracers to roll up. There is no mass diffusion in this problem, so that the two tracers should not interpenetrate along the interface. The results of CDG at the final time are shown in Fig. 2, and the method shows minimal artificial diffusion. Indeed, on this plotting scale, the CDG results are indistinguishable from the exact solution. An extreme example of an advection method that suffers from artificial diffusion is shown in Fig. 3. Here, the tracer quantities show significant diffusion along the interface. Such diffusion has been shown to have a large impact on the accurate prediction of geophysical flows [5].

For more information contact Robert Lowrie at lowrie@lanl.gov.

- [1] B. Cockburn, G.E. Karniadakis, C.-W. Shu, "Discontinuous Galerkin Methods: Theory, Computation and Applications," Lecture Notes in Computational Science and Engineering, Springer Verlag (2000).
- [2] W.H. Lipscomb, T.D. Ringler, *Mon. Weather Rev.* **133**, 2335 (2005).
- [3] B. van Leer, *J. Comp. Phys.* **23**, 276 (1977).
- [4] M.J. Prather, *J. Geophys. Res.* **91**, 6671 (1986).
- [5] S.M. Griffies, R.C. Pacanowski, R.W. Hallberg, *Mon. Weather Rev.* **128**, 538 (2000).
- [6] T. D. Ringler et al., *J. Comp. Phys.* (submitted).
- [7] The Parallel Ocean Program, Los Alamos Computer Code LA-CC 99-18.
- [8] M.A. Taylor, J. Edwards, A.S. Cyr, *J. Phys. Conf. Series* **125**, 12023 (2008).
- [9] M. Rancic, *Mon. Weather Rev.* **120**, 1394 (1992).

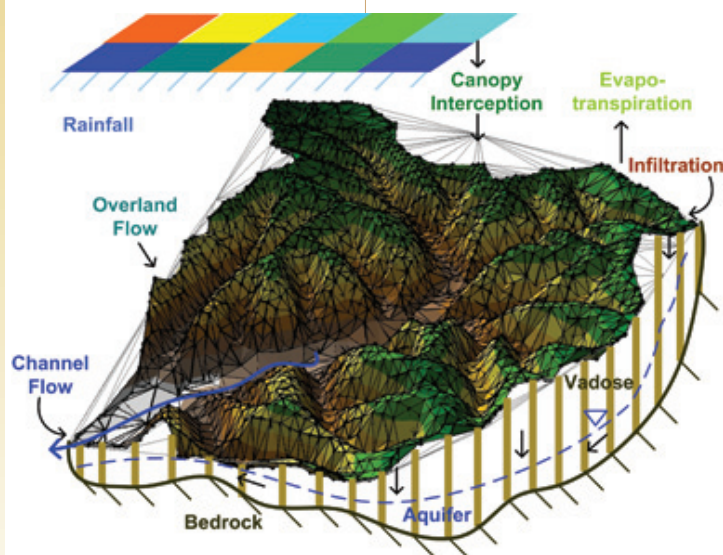
Funding Acknowledgments

DOE, Office of Science, Regional and Global Climate Modeling Programs

Subbasin-based Parallel Hydrological Models for Semi-arid Regions

Sue Mniszewski, Patricia Fasel, CCS-3; Everett Springer, STBPO-PRM;
Enrique Vivoni, Arizona State University; Amanda White, New Mexico Institute
of Mining and Technology

Fig. 1. The tRIBS model



A novel parallel approach has been applied to the TIN-based Real-Time Integrated Basin Simulator (tRIBS) [1,2], a physically based, distributed hydrologic model originally developed at the Massachusetts Institute of Technology. It is a C++ code that provides continuous simulation of watershed hydrology using an adaptive, multiple resolution representation of complex topography based on a triangulated irregular

High-resolution models of river basin hydrology can provide an important tool in assessing the impacts of global change, climate variability, and land-use changes on hydrologic response, leading to a more informed decision-making process about water resources and associated environmental issues. One of the primary obstacles towards advances in high-resolution watershed simulations has been the limited computational capacity available to most models. The growing trend of model complexity, data availability, and physical representation has not been matched by adequate developments in computational efficiency. This situation has created a serious bottleneck that has limited hydrologic models to small domains and short durations.

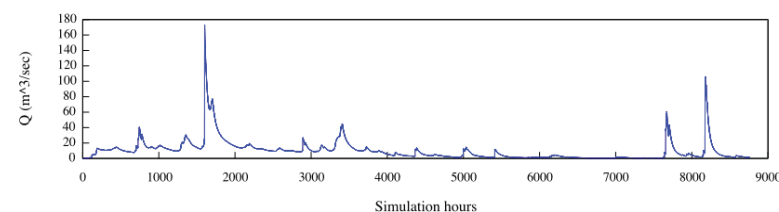


Fig. 2. Baron Fork hydrograph for November 1997–December 1998

network (TIN), modeling processes such as rainfall interception, evapotranspiration, moisture dynamics in the unsaturated and saturated zones, and runoff routing (see Fig. 1). In conjunction with the Surface Hydrology Research Group at Arizona State University and the NSF Science and Technology Center for Sustainability of Semi-Arid Hydrology and Riparian Areas (SAHRA) a parallel version of the tRIBS has been developed.

The parallel tRIBS approach [3] utilizes domain decomposition based on subbasins of a watershed. A stream reach graph of the channel network structure is used to determine each subbasin and its connectivity. Individual subbasins or subgraphs of subbasins are assigned to separate processors to carry out internal hydrologic computations (e.g., rainfall-runoff transformation). Routed stream-flow from each subbasin forms the major hydrologic data exchange along the stream reach graph. Individual subbasins also share subsurface hydrologic fluxes across adjacent boundaries. The parallel tRIBS code has been shown to run on high-resolution or large-sized basins (~1 million nodes).

In parallel tRIBS, as in the sequential version, hydrological processing occurs in three major steps: 1) unsaturated zone, 2) saturated zone, and 3) river routing. In the unsaturated zone step water flowing across the land surface and a short distance below surface is simulated. While all subbasins are being processed, flow data is sent from upstream ghost outlet nodes to downstream head nodes. After subbasin processing, state information from downstream head nodes is updated in the upstream ghost outlet nodes to be used in the next iteration. Next, in the saturated zone step, underground flow is simulated. The direction of flow

underground can change at each time step. Initial water table levels and ground water change are exchanged between processors that contain adjacent reaches. The river routing step simulates water flow through the stream reach network. Flow data from upstream ghost outlet nodes are sent to downstream head nodes during the calculation. Parallelization gains are achieved more from overlapping step processing than within step processing due to the dependencies between the related subbasins within the unsaturated zone and river network.

To improve usage of the parallel model, three supporting capabilities have been developed: 1) preprocessing construction of numerical meshes for large domains, 2) a mechanism to restart a long-running simulation, and 3) visualization for model debugging and results presentation. To speed up the mesh generation on all processors, a separate executable, MeshBuilder, is used to build input files containing the mesh, channel network, and ghost cell information specific to each processor. In the parallel operation, each processor is then able to read in and build only the assigned subbasins. Additional improvements related to faster calculations and sorting of the mesh generation also allow for the construction of very large domains ($>10^6$ nodes). Given time limitations in the use of cluster resources, a restart capability allows saving the entire model state at user-specified time intervals. This capability is useful for long (multiyear) simulations or for multiple scenarios from a given state. Finally, a tRIBSReader Visualizer plugin library for the ParaView visualization toolkit [4] allows for the rapid display, inspection, and animation of the static basin characteristics and dynamic model output.

An efficient partitioning balances the computational load and minimizes message passing between processors. The number of nodes per reach contributes to the computational load. Connections between reaches in the stream network and subsurface flux network contribute to the messaging. Four methods were considered to improve load balancing. The default method slices up the reaches as created by tRIBS. Three others were produced using METIS multiconstraint graph partitioning. The flow method balances

the number of nodes (computational load) and stream reach network connectivity. The flow-flux method additionally considers the subsurface flux connectivity. And the flow-flux-upstream method also balances the number of reaches without upstream reaches.

A high-resolution Oklahoma Baron Fork basin model (~900K nodes) was run for a 1-year period from November 1997 to December 1998 (see Fig. 2) for performance and load balancing experiments on the Institutional Computing Coyote cluster. The default method randomly distributes reaches across processors (see Fig. 3), does not balance the computational load, and results in excessive message exchanges, especially for the subsurface flux. The flow, flow-flux, and flow-flux-upstream methods all produce geographic distributions of subbasins based on the flow network (see Fig. 3). The computational load as nodes is distributed more evenly for these methods, while flow minimizes the river routing exchanges and flow-flux minimizes the subsurface flux exchanges between processors. Run times are greatly improved compared with the default method by up to a factor of 10 for 32, 64, 128, and 256 processors. The best performance is seen for flow-flux on 128 processors with 2.186 min/simulated day. Compared with run time per simulated day on one processor, this method results in a 42X speedup.

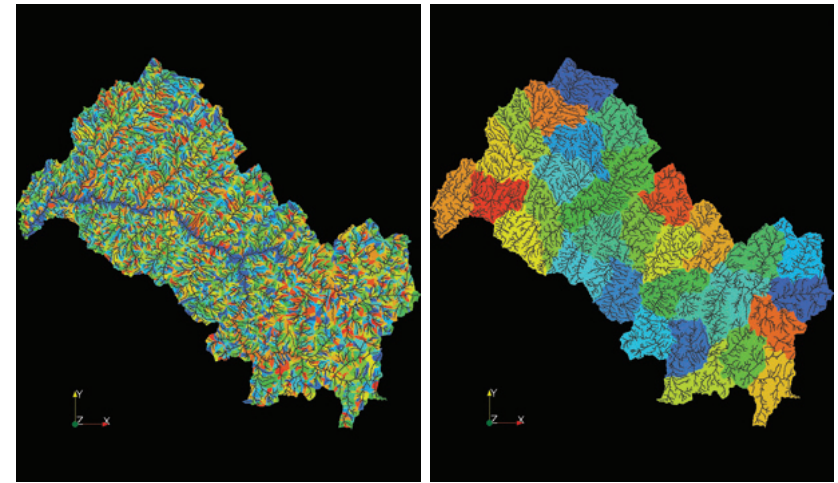


Fig. 3. Baron Fork sub-basins colored by partition on 32 processors for default (left) and flow-flux (right) load balancing methods

**For more information contact
Sue Mniszewski at
smm@lanl.gov.**

- [1] V.Y. Ivanov et al., *Water Resour. Res.* **40**, W11102 (2004).
- [2] E.R. Vivoni et al., *Hydrol. Earth Sys. Sci.* **11**, 1683 (2007).
- [3] S. Mniszewski et al., *Increased Efficiency for High-Resolution Baron Fork Simulations Using Basin Structure Characteristics*, 9th SAHRA Annual Meeting (2009).
- [4] J. Ahrens et al., *ParaView: An End-User Tool for Large-Data Visualization in The Visualization Handbook*, Academic Press, 717 (2005).

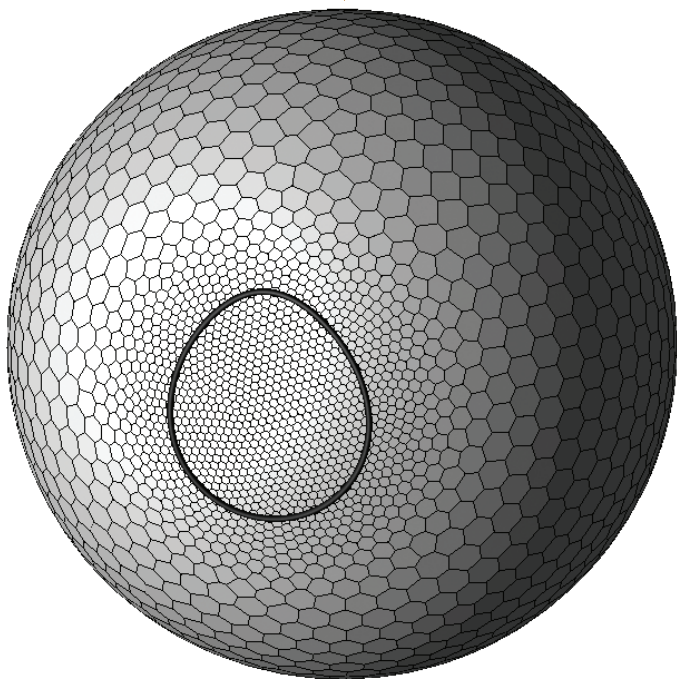
Funding Acknowledgments

LANL Directed Research and Development Program

Bridging Scales in the Earth's Climate System through Multiresolution Modeling

Todd D. Ringler, T-3

Fig. 1. An example spherical centroidal Voronoi diagram [3]. In this simple example, the mesh is enhanced in a specific region of interest, in this case a topographical feature. Lower resolution is used elsewhere.



In the coming decade and beyond, the climate modeling community will be challenged to resolve scales and processes that are far beyond its current scope. The challenge to resolve fine-scale and new processes will arise primarily for two reasons. First, unresolved processes that have a significant influence on the global climate system are likely to remain. Examples of such unresolved processes might include ice streams within large-scale ice sheets, ice shelves collapsing due to interaction with ocean processes, cloud processes in the atmosphere, and the dependence of ocean biochemistry on ocean

eddy activity. Second, there is the pressing need to quantitatively characterize the regional-scale signature of anthropogenic climate change. A regional example is the climate-change driven impacts to hydrological processes in the western United States.

The ability to resolve the scales noted above throughout the relevant climate system component is beyond the current computing capacity of even the most powerful computers available today. Basic scaling arguments of computer resources available into the future indicate that this situation will remain

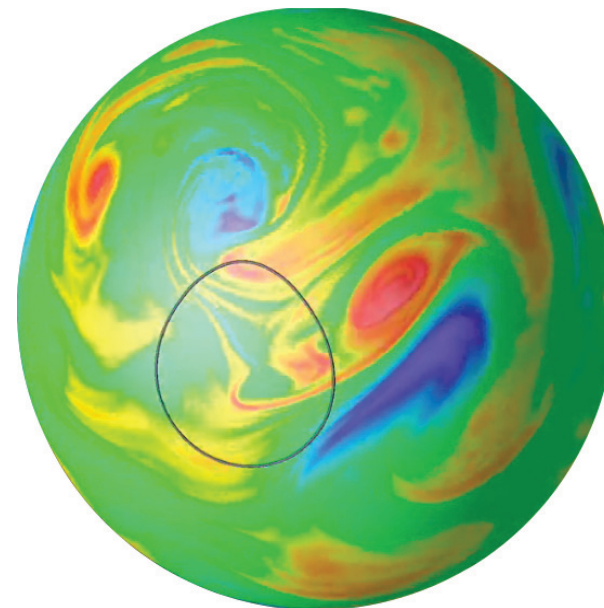


Fig. 2. Using a higher resolution version of the mesh shown in Fig. 1, a 2D turbulence simulation is conducted where the forcing is generated by an isolated topographic feature around which the mesh is enhanced. The field shown is relative vorticity. The simulation remains smooth even as strong filamentation occurs.

for decades to come. As a result, the climate modeling community is faced with the challenge of creating numerical algorithms able to support multiple resolutions within a single global climate model simulation, where relevant processes are resolved in specific regions as opposed to being resolved globally. The creation of what is effectively a multiscale climate system model will allow the scientific community to explore the impacts of currently unresolved processes in specific geographic regions. Such a model will be of critical use to both society and policy makers in order to better understand the regional impacts of global climate change.

While it is clear that access to a multiscale climate system model would be of great benefit to both science and society, several fundamental challenges have precluded the creation of such a model to date. The first challenge is the creation of numerical algorithms

that can produce robust, stable simulations over century time-scale integrations when the resolution of the underlying grid varies substantially from one location to another. The second challenge is related to the physical parameterizations included in climate system models. In a climate model that supports multiple resolutions, physical parameterizations are required in low-resolution areas where direct simulation is not possible. Yet these same physical parameterizations must be excluded in high-resolution regions where the model is of sufficient resolution to directly simulate the phenomena. The creation of scale-aware physical parameterizations needed for climate system modeling has been elusive.

Along with our partners at the National Center for Atmospheric Research (NCAR) and Exeter University, the climate modeling group at LANL has made significant progress overcoming the first challenge listed above. Essentially, we have derived a finite-volume method that mimics the relevant aspects of the underlying equations of fluid motion, even when the method is situated on a mesh with multiple resolutions [1,2]. While the numerical algorithm is applicable to a wide class of meshes, it is most naturally paired with a spherical centroidal Voronoi diagram [3]. Figure 1 shows an example of a spherical centroidal Voronoi diagram where the surface of the sphere is decomposed into two dominate resolutions, with low-resolution covering most of the Earth and high-resolution used in the area of interest. When we use the multiresolution finite-volume method along with the mesh shown in Fig. 1, we find that we can simulate 2D turbulence without the need for ad hoc stabilization methods (see Fig. 2). In 2D turbulence, energy cascades upscale while potential enstrophy cascades downscale. Remarkably, the variable mesh simulation confirms our analytical findings that the discrete model conserves total energy to within time-truncation error, thus producing an essentially inviscid simulation. Over the next year the climate modeling group at LANL will create global atmosphere and ocean climate model components that can utilize this variable-mesh technology. The first steps toward this goal

have already been realized. Figure 3 shows a snapshot of potential vorticity for an eddy-resolving ocean basin simulation.

While a tremendous amount of effort is required before the full breadth of this approach can be realized, the idea has the potential to fundamentally change how global climate system models are used and developed. By allowing regions of local mesh refinement, our approach will enable new lines of scientific inquiry focused on climate processes with 1 km to 25 km length scales to begin decades before such resolutions will be resolved in traditional, globally-uniform climate system models. In essence, the approach will hopefully create a framework for the simulation of innovative, fine-scale climate processes within the context of global climate system modeling.

For more information contact Todd Ringler at ringler@lanl.gov.

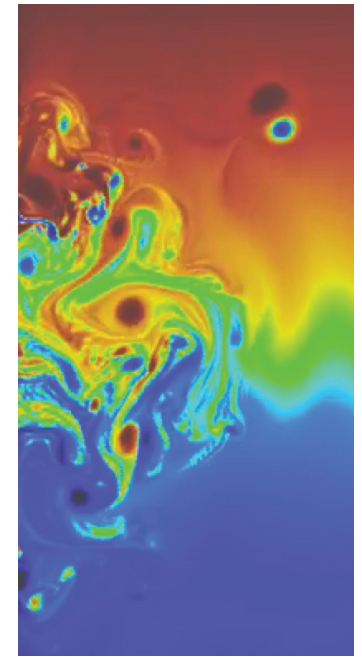


Fig. 3. A snapshot of potential vorticity from an idealized eddy-resolving ocean basin simulation.

- [1] J. Thuburn et al., *J. Comput. Phys.* **228**, 8321 (2009). doi:10.1016/j.jcp.2009.08.006.
- [2] T. Ringler et al., *J. Comput. Phys.*, (2010). 10.1016/j.jcp.2009.12.007.
- [3] T. Ringler, L. Ju, M. Gunzburger, *Ocean Dyn.* **58**, 475 (2008). doi:10.1007/s10236-008-0157-2.

Funding Acknowledgments

DOE Office of Science, Regional and Global Climate Modeling Program

Information Science and Technology

The Information Science and Technology field draws upon computer science, computational science and mathematics to provide innovation in the areas of data processing and management, and deriving actionable information from data. Information Science and Technology capabilities underly most science endeavors at LANL.

This section includes contributions to the field of Information Science and Technology in areas ranging from applied computer science to applications for decision support. Three articles in this section address the vital areas of efficient data storage, data archiving, and resilient computing in applied computer science research and development. The remaining four articles address methods for harnessing Information Science and Technology to apply to real-world data-intensive problems in the areas of portal monitoring, image reconstruction, uncertainty quantification for weapons management and cyber threat analysis.

Information Science and Technology



PLFS: A Checkpoint Filesystem for Parallel Applications

John Bent, HPC-5; Garth Gibson, Carnegie Mellon University; Gary Grider, HPC-DO;
Ben McClelland, HPC-3; Paul Nowoczynski, Pittsburgh Supercomputing Center;
James Nunez, Milo Polte, Meghan Wingate, HPC-5

Given the scale of massively parallel systems, system failures are not a matter of if, but of when. When an inevitable system failure takes place, it is not viable to restart the calculation from the beginning, particularly if such an application (e.g., a complex simulation) has been running for a month or so. Thus, massively parallel applications rely on checkpointing, which consists of saving snapshots of the current state of an application into files on persistent parallel storage systems. After a failure, users can restart the application from the most recent snapshot.

At present, concurrent and random-access checkpoint writes to a shared file require unnecessary 1) disk-seeking, and 2) file-locking at the parallel storage system. Such requirements drastically reduce bandwidth by as much as two orders of magnitude. To overcome this problem, we developed software known as Parallel Log-structured File System (PLFS), which decouples concurrent access and reorganizes logical writes into sequential physical writes, thus enabling applications to write in parallel at a near-optimal storage bandwidth.

Fig. 1. How PLFS rewrites a checkpoint file.

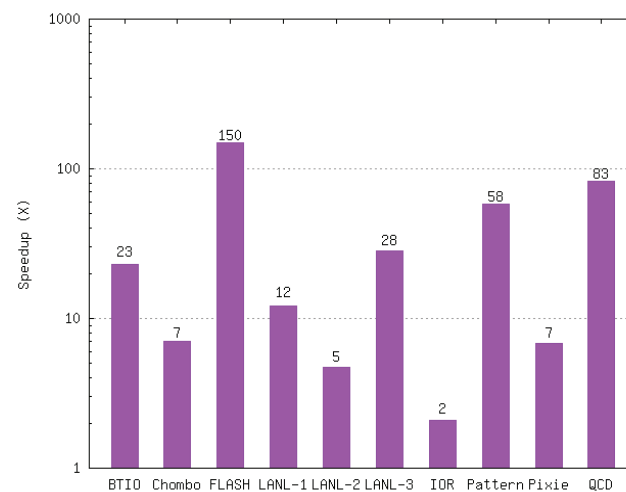
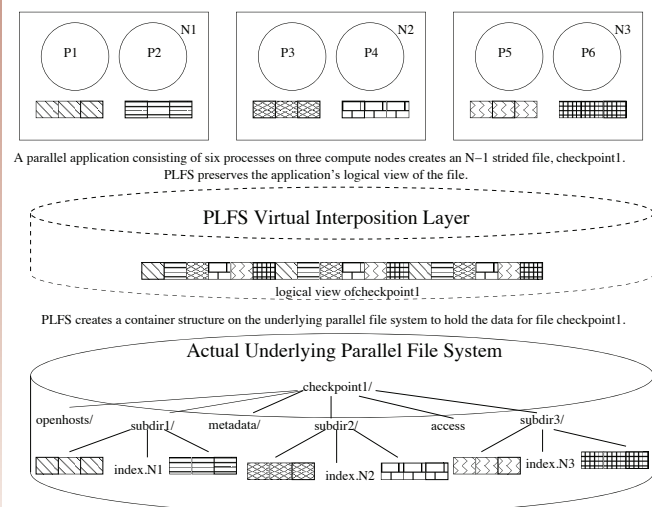


Fig. 2. These bars show the speedup in write bandwidth for eight different applications and two synthetic benchmarks. Note that the y-axis is logscale.

PLFS uses the data and metadata within the container to implement a virtual interposition, or intermediate, layer. Multiple processes opening the same logical file for writing share the container, although each opening receives a unique data file within the container into which all of its writes are appended. By giving each writing process in a parallel application access to a nonshared data file, PLFS transparently converts write-access patterns to improve checkpoint bandwidth by as much as several orders of magnitude.

Figure 1 shows how PLFS reorganizes a N-1 checkpoint file onto an underlying parallel system. The processes create a new file on PLFS called checkpoint1, causing PLFS in turn to create a contained structure on the underlying parallel file system. The figure also shows the access file, which is used to store ownership and privileged information about the logical file, the openhosts, and the metadata directories, all of which are used to cache metadata to improve query time.

File systems have two basic checkpointing patterns: N-N and N-1. An N-N checkpoint is one in which each of a certain number of processes (N) writes to a unique file, for an equal number of files written (N processes = N files written). An N-1 checkpoint differs

in that all N processes write to a single shared file. Although both N - N and N -1 patterns pose challenges, we have observed that the challenges of N -1 checkpointing, from the perspective of the parallel storage system, are much more difficult. Applications using N -1 patterns consistently achieve significantly less bandwidth than do those using an N - N pattern. Because N - N checkpointing derives higher bandwidth than N -1, the path to faster checkpointing is for application developers to rewrite existing N -1 checkpointing applications to perform N - N checkpointing instead. Additionally, all new applications should be written to take advantage of the higher bandwidth available to N - N checkpointing. Some developers have gone to N - N checkpointing, but many continue to prefer an N -1 pattern even though its disadvantages are well understood.

Of the 23 applications listed on the Parallel I/O Benchmarks, an industry standard developed by the Parallel I/O Benchmarking Consortium, at least 10 have an N -1 pattern. Two major applications at LANL use an N -1 checkpointing pattern, as do at least two of the eight initial applications chosen to run the Laboratory's Roadrunner, the world's first petascale computer. The parallel storage system attached to Roadrunner is the largest the Laboratory has ever had; testing it has revealed that the challenges of N -1 patterns are severely exacerbated at this scale. Given current bandwidths, we know of no current N -1 application at LANL that can effectively checkpoint across the full width of Roadrunner. PLFS, currently mounted on Roadrunner and many other Laboratory systems, allows them to do so (Fig 2). More generally, PLFS achieves this improvement for any parallel writes and not just checkpoint writes. PLFS also has been shown to improve bandwidth for many parallel reads.

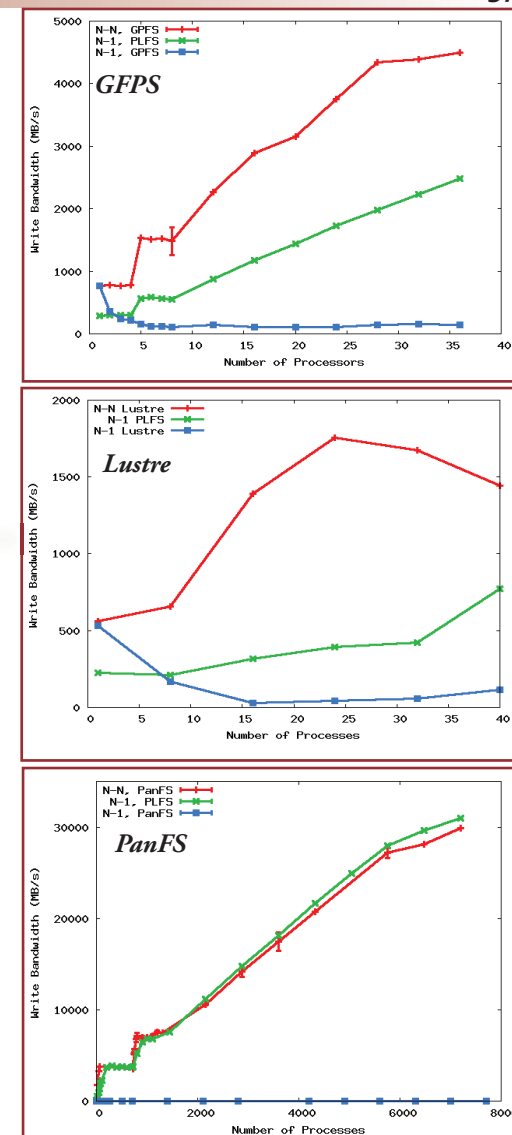
The theory underlying this work is that concurrent, random-access checkpoint writes to a shared file require unnecessary disk-seeking and file locking at the parallel storage system, thereby drastically reducing bandwidth. By decoupling concurrent access and by reorganizing random logical writes into sequential physical writes, PLFS enables applications to write in parallel at a near-optimal storage bandwidth.

In addition to observing a large performance improvement at the Laboratory, which uses the Panasas ActiveScale (PanFS) storage system, we have also tested PLFS on the two other most widely used storage systems for supercomputing: the Linux cluster storage system (Lustre) and GPFS, the IBM General Parallel File system. We have demonstrated a 25X speedup on GPFS, an 8X speedup on Lustre, and a 3000X speedup on PanFS. Notice that the Lustre and GPFS experiments were run on very small clusters. Extrapolating the trends and comparing with the PanFS results suggests that speedups for GPFS and Lustre on reasonably sized clusters will be similar to the PanFS results.

Figure 3 shows the impact of PLFS on three different parallel storage systems: GPFS, Lustre, and Panasas (PanFS). In each graph, the y-axis shows the write bandwidth controlled by the number of concurrent writers on the x-axis. The red lines show the bandwidth for N - N workloads, and the blue lines are the bandwidth for small, strided N -1 workloads (*strided* being a type of regular access pattern).

Notice that all three storage systems show massive performance degradation for these N -1 workloads. The green lines are the performance of these same N -1 workloads as written through PLFS. PLFS allows the application to use its preferred N -1 workload and achieve close to the bandwidth it would have achieved if using an N - N workload. For Lustre and GPFS, the N -1 bandwidth through PLFS is much better than the N -1 bandwidth directly to the storage system, but it is not as good as the N - N bandwidth.

For Panasas, however, N -1 through PLFS does match the N - N bandwidth. Notice that the x-axes are very different. The GPFS and Lustre results were gathered on small clusters of fewer than 50 nodes, whereas the Panasas results come from a cluster of several thousand nodes. We expect that the PLFS results on Lustre and GPFS will match N - N bandwidths for larger-sized clusters.



For more information contact
John Bent at
johnbent@lanl.gov

Funding Acknowledgments

- National Science Foundation (NSF)
- DOE, NNSA Advanced Simulation and Computing (ASC) Program
- DOE, Scientific Discovery through Advanced Computing Program (SciDAC)

Data Mining in Radiation Portal Monitoring

Tom Burr, Michael Hamada, CCS-6; Nicolas Hengartner, CCS-2; Kary Myers, Richard Picard, CCS-6

Currently deployed passive gamma and neutron detectors screen for illicit nuclear material. Archived data can help to evaluate special nuclear material detection probabilities (DP) and to investigate several related issues, including: 1) nuisance gamma alarms arising from naturally occurring radioactive materials (NORM), 2) sensor fusion options, and 3) radioisotope identification (RIID) performance.

Figure 1 shows a screening location where four detector panels each record a neutron and a low- and high-energy gamma count every 0.1 s for 5–20 s, resulting in a 12-component time series of 50 to 200 observations.

Nuisance alarms due to NORM reduce DPs for threats. Strategies for recognizing common NORMs, such as cat litter or ceramics, depend on the sensor energy resolution. Figure 2 illustrates the extent to which different, common NORM categories have a signature using a 2D representation for each profile. The scaled maximum low-energy count range vs the scaled maximum ratio of the high- to low-energy counts is plotted for each of several profiles of eight NORM categories. One of the best methods using the systems described here (two-energy gamma and neutron in each of four panels) uses a nonparametric density-estimation method for pattern recognition [1,2] applied to tens of features such as those used in Fig. 2, derived from the 12-component time series for each profile. Although some common NORMs do appear to have a signature, at present any vehicle having a large count is subject to further investigation. Such additional investigation and measurement result in slower vehicle transit times.

Vehicles that alarm in primary screening go to secondary screening where higher resolution gamma and X-ray measurements are made. There are several feasible options for combining information from primary and secondary screenings to enhance DPs. Some options have been quantitatively evaluated [3,4]. By sensor optimization, we mean to optimize the expected DPs with respect to the sensor thresholds, sensor ordering and protocol, and/or the alarm rules. Sensor protocol involves, for example, whether the sensors send only pass-fail information or more complete data, and whether the declared cargo is allowed to impact the calculated likelihoods.

RIID performance is a major challenge for low-, medium-, or even high-resolution gamma spectra. On the basis of a small test data set, we have found that medium-resolution detectors such as the hand-held sodium iodide (NaI) detectors used in secondary screening appear to be competitive with high-resolution detectors. A current challenge is to evaluate the cost/benefit that medium-resolution NaI detectors might provide in primary screening, deployed as so-called advanced spectroscopic portals (ASP). ASP testing is ongoing to estimate low-, medium-, and high-resolution detector performance based on several metrics. One straightforward opportunity to improve RIID performance and testing appears to be spectral smoothing; adjustments are made to preserve key spectral regions of interest such as peak areas. RIID algorithms can be tested for more measurement scenarios by reducing count time and by augmenting real spectra with realistic synthetic spectra. Model uncertainty will play a key role in assessing the adequacy of synthetic spectra.

For more information contact Tom Burr at tburr@lanl.gov.



Fig. 1. Example screening location with four detector panels surrounding the vehicle.

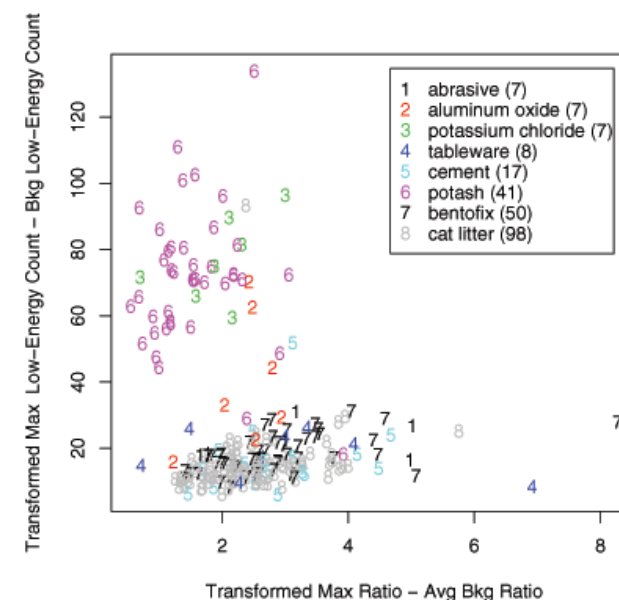


Fig. 2. The scaled maximum low-energy gamma count versus the scaled maximum ratio (adjusted by subtracting the corresponding average values in the background data taken just prior to the vehicle profile) for each of several profiles of eight common NORMs. The number of profiles for each NORM category is given in parentheses in the figure legend.

- [1] T. Burr, K. Myers, *Appl. Radiat. Isot.* **66**, 1250 (2008).
- [2] T. Burr, J. Doak, *Intell. Data Anal.* **11**, 651 (2007).
- [3] T. Burr et al., "Sensor Optimization for Radiation Portal Monitoring," *Proc. Inst. Nucl. Mater. (Management Annual Meeting)* (2006).
- [4] R. Picard, T. Burr, *IEEE Sens. J.* **8**, 1389 (2008).
- [5] T. Burr et al., *Qual. Reliab. Eng. Int.* **25**, 899 (2009).
- [6] T. Burr, M. Hamada, *Algorithms* **2**, 339 (2009).
- [7] T. Burr et al., "Smoothing Low-Resolution Spectral Data," *IEEE Trans. Nucl. Sci.* (submitted).

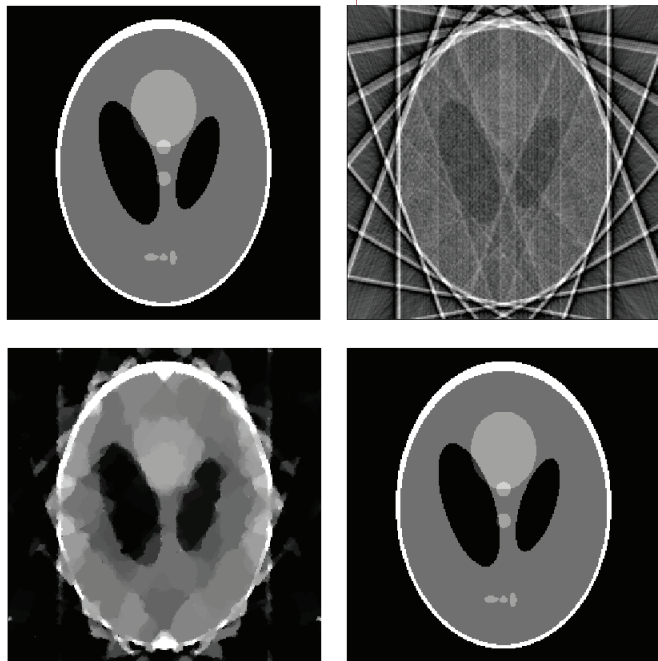
Funding Acknowledgments

Department of Homeland Security (DHS)

Fast Algorithms for Image Reconstruction from Very Few Data

Rick Chartrand, T-5

Fig. 1. Top left: the Shepp-Logan phantom, a test image. Top right: reconstruction using filtered back-projection is poor. Bottom left: reconstruction using usual, convex compressive sensing method is also poor. Bottom right: our nonconvex reconstruction is exact.



Recent developments in applied mathematics have demonstrated that images and other signals can be reconstructed from far fewer data than traditionally believed necessary [1,2]. The results exploit the fact that images of the real world or human-generated experiments are a very special subset of all possible images, the overwhelming majority of which will look like random noise. When we try to reconstruct an image from data, what sets apart the image we are looking for is that it is sparse – it can be represented with relatively few coefficients, far fewer than the number of pixels in the image. We have developed reconstruction algorithms that can reconstruct images from a number of data that is commensurate

with the complexity (or information content) of the image, rather than its size. These algorithms are at the forefront of the new and popular field known as compressive sensing, so called because the few data from which an image can be effectively reconstructed constitute a compression of an image. Instead of needing to obtain every pixel of an image and then compressing it (as is done by digital cameras, for example), we can, in effect, measure a compression directly. There are many applications of the ability to reconstruct images from very few data, particularly whenever data are difficult, expensive, or dangerous to acquire.

For example, in medical imaging the high X-ray dose of a CT scan can be replaced with relatively few radiographs. Better yet, in many cases the much safer MRI can be substituted—the usual barrier is the high cost of the procedure, due to the long scanning time required. Compressive sensing techniques can allow a much shorter scan to be used instead, making the procedure more palatable to insurance companies. Both CT and MRI are also used in national security applications, where the ability to reduce data collection time can be very advantageous. Many other applications related to LANL's mission stand to benefit as well. In remote sensing or space situational awareness, one is often fundamentally limited in the amount of information one can gather. Compressive sensing would allow more information to be extracted from the data that are available.

As a simple example, consider Fig. 1. We reconstruct a test image, the Shepp-Logan phantom, which was designed to be challenging for medical imaging algorithms to reconstruct. Our data are samples of the 2D Fourier transform of the phantom, taken along nine radial lines through the origin of the frequency domain, together making up less than 3.5% of the full Fourier transform. Sampling along radial lines in this way makes the data roughly equivalent to having radiographs of the phantom, one for each line. The reconstruction problem can thus be seen as a limited-view CT problem, but also as an MRI problem, as MRI data can be seen as directly sampling the Fourier transform as the object. Our reconstruction takes advantage of the fact that the gradient of the phantom is very sparse, being zero except at the boundaries of the ellipses. Our reconstruction approach is to solve an optimization problem, which minimizes a sparsity-inducing penalty term, subject to the data constraint. The usual reconstruction method for CT is filtered backprojection, which gives a very poor reconstruction, being designed for having hundreds of radiographs. The usual compressive sensing approach also fails with so few data, but our particular method gives an exact reconstruction.

The reason our compressive sensing approach outperforms the usual one is that it is a closer approximation to what we really want to solve, namely the problem of simply finding the sparsest solution that is consistent with the data. However, directly solving that

problem is computationally intractable; all the world's computers working for trillions of years would hardly make a dent. The field of compressive sensing was born when it was discovered that one could replace this problem with one that is convex. Convexity of a function implies that any local minimum will be a global minimum—this means, for example, that one can just head downhill, such as moving in the direction opposite of the gradient of the function, and eventually get to a global minimum. The research field of convex optimization is very mature, so there are several computationally efficient algorithms available for solving the convex approximation of the sparse recovery problem. Under reasonable conditions, it has been proven that the solution of the convex problem will be the same as that of the intractable sparse recovery problem. This result generated substantial excitement, with compressive sensing now being one of the fastest growing areas of applied mathematics.

Our approach, however, is nonconvex. Simply heading downhill will almost certainly result in convergence to a local minimum, but not the global minimum. For the Shepp-Logan phantom example, the number of local minima exceeds $10^{4,500}$! This makes it seemingly impossible to hope to obtain the global minimum. However, we have developed an algorithm approach that does so, with tremendous reliability [3,4]. The key is an iterative smoothing approach, which first smooths out the function being minimized so that it no longer has any local minima. The smoothing is then successively diminished, with the solution at each stage used to initialize the next stage. The result is that the iteration reaches the right convergence basin, before the local minima reappear. We thus obtain algorithms that are able to avoid the local minimum problem, and also serve to more closely approximate the true sparse recovery problem. This is what allows our algorithms to successfully reconstruct images from many fewer data than the usual, convex approach to compressive sensing.

Our latest research has led to an algorithm that not only works with fewer data than its predecessors, but is computationally very fast, at least for many image reconstruction problems [5]. It works

particularly well for cases where the data consist of samples of the Fourier transform, as in the example above. In addition to CT and MRI, Fourier-domain sampling arises in applications such as synthetic aperture radar and sonar, or where interferometry is used, such as in radio-astronomy or spectroscopy. Most algorithms that converge in a reasonable number of iterations require solving a large, linear system during each iteration. These can be computationally expensive to solve, especially for large-scale problems. Our approach allows this system to be solved directly in the Fourier domain, essentially requiring just two Fast Fourier Transforms (FFTs). Using FFTs makes solving the linear system much faster, while also scaling very well to large problems. For example, the first Shepp-Logan phantom experiments took the better part of an hour to reconstruct the 256×256 image. The new method takes just seconds, using very simple prototype code (in Matlab) on a simple laptop (see Fig.2). Thus, using nonconvex optimization, with unparalleled abilities to recover images from fewer data than ever before, image reconstruction can now be done using algorithms that are efficient, making many large-scale applications much more feasible.

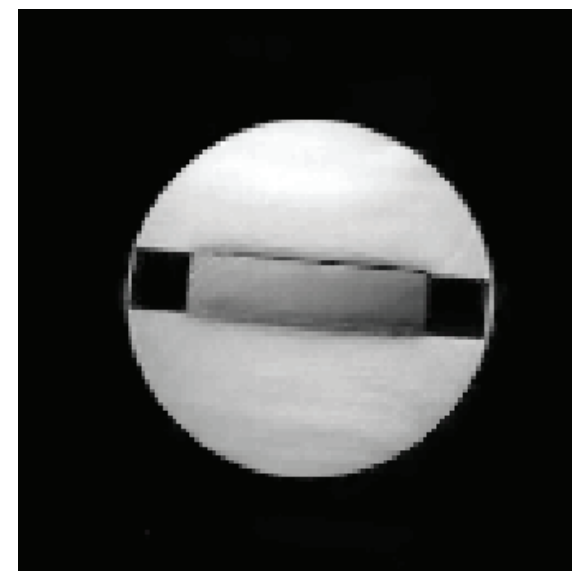


Fig. 2. Reconstruction of a real test object from 8-coil MRI data with 31% sampling. All eight reconstructions were performed in 17 s, using simple Matlab code on a laptop.

**For more information contact
Rick Chartrand at
rickc@lanl.gov.**

- [1] E. Candès, J. Romberg, T. Tao, *IEEE Trans. Inf. Theory* **52**, 489 (2006).
- [2] D.L. Donoho, *IEEE Trans. Inf. Theory* **52**, 1289 (2006).
- [3] R. Chartrand, *IEEE Signal Processing Lett.* **14**, 707 (2007).
- [4] R. Chartrand, W. Yin, *IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, (2008).
- [5] R. Chartrand, *IEEE Int. Symp. Biomedical Imaging*, (2009).

Funding Acknowledgments

- LANL Directed Research and Development Program
- Defence Advanced Research Projects Agency (DARPA)—Multi-aperture Geosynchronous Imaging Project

COTS Parallel Archive Integration Experiences

Gary Grider, HPC-DO; Hsing-Bung Chen, HPC-5; Cody Scott, Milton Turley,
Aaron Torres, Kathryn Sanchez, John Bremer, HPC-3

There is demonstrated need for current and future archive storage systems for high performance computing to 1) scale to very high bandwidths, 2) scale in metadata performance, 3) support policy-based hierarchical storage management capability, 4) scale in supporting changing needs of very large data sets, 5) support a standard interface, and 6) use commercial-off-the-shelf (COTS) hardware. This is akin to changes in parallel file systems, although file systems perform at one or more orders of magnitude faster. Archive systems continue to move closer to file systems in their design due to the need for speed and bandwidth, especially metadata searching speeds, by using more caching and less robust semantics. Currently the number of extreme highly scalable parallel archive solutions is very small, especially those that will move a single large highly striped parallel disk file to many tapes in parallel. We are developing a hybrid storage approach, using COTS components and innovative software technology to bring new capabilities to productive use for the HPC community much faster than if we create and maintain a complete end-to-end unique parallel archive software solution.

In this project we worked to integrate a COTS global parallel file system and a standard backup/archive product with a very small amount of additional user space code to provide a scalable and parallel solution that overlaps highly in function with current niche parallel archive product(s), including 1) doing parallel movement to/from tape for a single large parallel file, 2) hierarchical storage management, 3) Information Lifecycle Management (ILM) features, 4) high volume (nonsingle parallel file) archives for backup/archive/content management, and 5) leveraging all free file movement/management tools in Linux such as copy, move, ls, or tar.

The Advanced Simulation and Computing (ASC) program at LANL has a goal to pilot a more COTS-based archive in less than 5 years. We had an opportunity to do a pilot archive project for the LANL Roadrunner cluster while it was in its testing phase. We chose the IBM General Parallel File system (GPFS) for the parallel file system because of the new ILM features. We chose the IBM Tivoli Storage Manager (TSM) because we had it in-house. We have designed, developed, and integrated the following features in the proposed COTS Parallel Archive System (Fig. 1, Fig.2):

- A parallel tree walker and copy user space utility
- A storage pool (stgpool) support in utility (using file system application program interface (API))
- An efficient ordered retrieval in utility (using dmapi API and back-end tape system query)
- Archive Parallel File System to support ILM stgpool features
- Archive back-end to support ILM stgpool and colocation features
- A Filesystem in UserSpace (FUSE) to break up enormous files into parts that can be migrated and recalled in parallel to/from back-end tape system

We show our proposed COTS Parallel Archive System in two parts: the back-end system (Fig. 3) and the front-end system (Fig. 4a and b).

We have built and applied our working COTS Parallel Archive System to one of the world's fastest computer systems, Roadrunner. We have successfully supported the seven Roadrunner Open Science Projects to archive over 1 petabyte of data and demonstrated its capability to address requirements of future archive storage systems. This COTS system is now providing archive services in LANL's Turquoise open collaboration network.

For more information contact Cody Scott at cscott@lanl.gov.

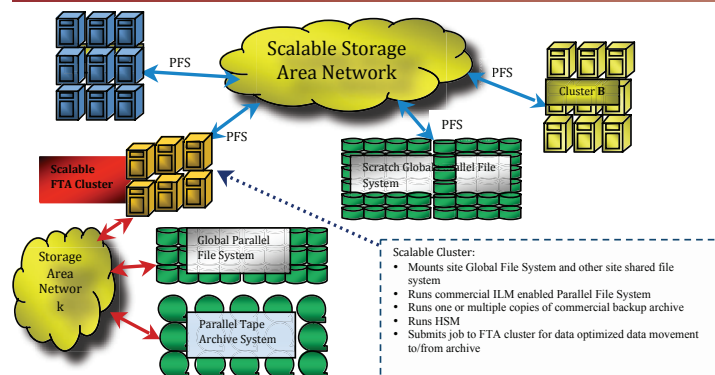


Fig. 1. Proposed COTS parallel archive system to deploy a parallel archive with a parallel file system as its first tier of storage.

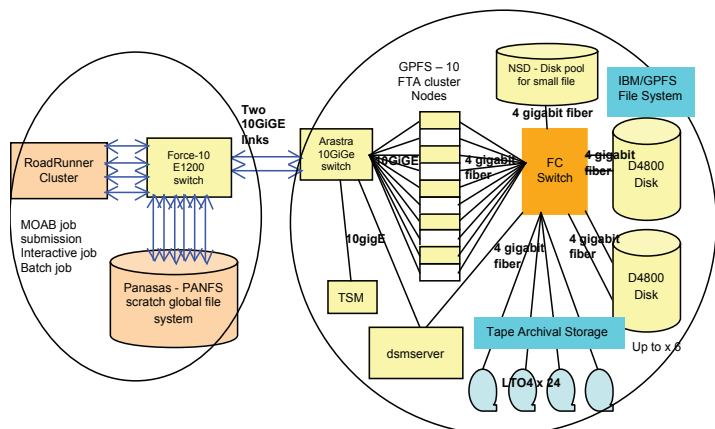


Fig. 2. COTS parallel archive system for LANL's Open Science Projects.

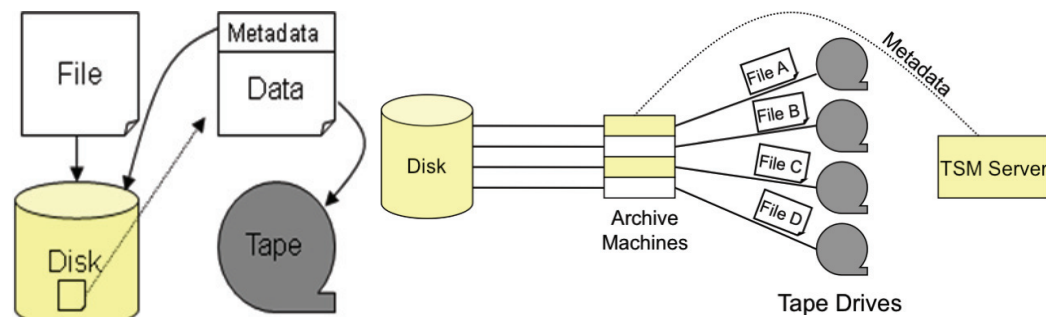


Fig. 3. The backend system of the COTS parallel archive system.

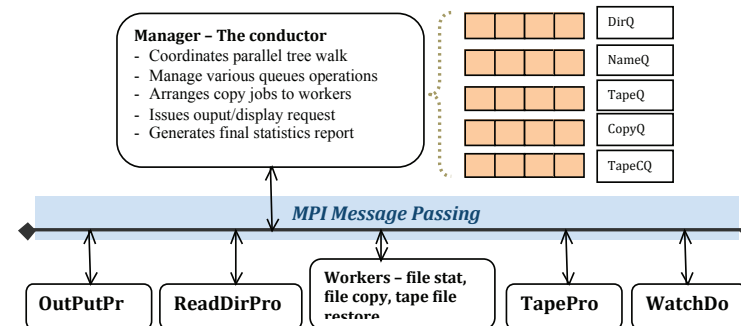


Fig. 4a. PFTool software system diagram.

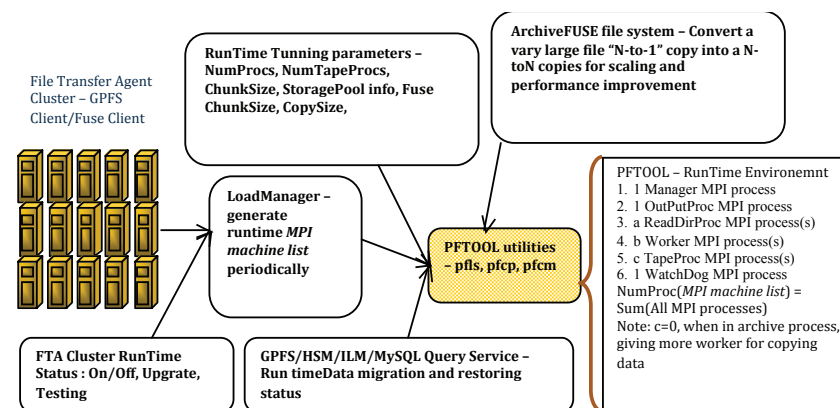


Fig. 4b. PFTOOL system runtime operation diagram.

Funding Acknowledgments

LANL Institutional Computing Program

Report on Current Issues in Resilience

Nathan DeBardeleben, HPC-5

High-end computing (HEC) is requisite for solving our nation's most important scientific and engineering problems and has become increasingly vital to the mission of the national security community [1]. As the scale and complexity of HEC systems continue to grow, the impact of faults and failures will make it increasingly difficult to accomplish productive work using traditional means of fault-tolerance [2,3]. Further, the challenges of integrating large complex heterogeneous systems are increasing to the point where the stabilization period consumes a significant portion of the lifetime of those systems [4,5]. As a consequence of these two troubling trends, it will be necessary for the HEC community to identify innovative means for efficiently and affordably performing productive work on systems encountering frequent, persistent, and erratic errors—many of which will be undetectable by existing system-monitoring solutions. Resilience meets these daunting and ever-increasing challenges. To ensure the continued viability of the largest, most powerful, leading-edge computing systems will require standards-based solutions. These solutions must efficiently and dynamically guard and preserve information, computation, and data movement in the presence of faults and failures arising from complex system interactions and dependencies among platform hardware and software components, the system workload, and the physical environment.

The goal of HEC resilience is to enable effective and resource-efficient use of computing systems at extreme scale in the presence of system degradations and failures.

At the highest level, high-end computing is the process by which data is transformed into information through computation. Resilience facilitates this critical transformation process by accepting that the underlying hardware and software that comprises a system will be unreliable. In order to succeed, resilience assumes a new perspective in which uncertainty about the state of the system plays an important role in managing that system. Resources traditionally focused on maintaining a known and desirable system state are instead focused on end-to-end fidelity of data, computation, and data movement. Resilience is concerned with reliability of information in lieu of, or even at the expense of, reliability of the system. This novel approach to fault-tolerance is necessary in order to address the two-fold challenge of decreasing system reliability, because of increasing scale, and decreasing certainty about the operational state of the system due to increasing complexity. The resilience community proposes to address these challenges in five focused but overlapping thrust areas: (see Fig.) 1) theoretical foundations, 2) enabling infrastructure, 3) fault prediction and detection, 4) monitoring and control, and 5) end-to-end data integrity. To manage this prodigious scope, a successful program of resilience research will require coordinated, multidisciplinary undertakings in each of these thrust areas. This requirement forms the justification of a call for a national effort in resilience.

For more information contact Nathan DeBardeleben at ndebard@lanl.gov.

- [1] Report of the *High End Computing Revitalization Task Force* (HECRTF), May (2004).
- [2] B. Shroeder, G.A. Gibson, *J. Phys.: Proc. Sci. Disc. Adv. Comput. Prog. (SciDAC) Conf.* **78**, 2022 (2007); <http://www.iop.org/EJ/abstract/1742-6596/78/1/012022>.
- [3] J.T. Daly, "Application Resilience for Truculent Systems," at *Fault Tolerance Workshop for Extreme Scale Computing* (2009); <http://www.teragridforum.org/mediawiki/images/8/80/Daly2009ws.pdf>.
- [4] "Roadrunner to Expand Hybrid Computing Applications," *ASC eNews Quarterly News Letter*, NA-ASC-500-09 Issue 10.
- [5] B. Comes, B. Bland, "Testing & Integration," *Petascale Systems Integration into Large Scale Facilities Workshop* (2007); http://www.nersc.gov/projects/HPC-Integration/presentations/Breakout_3_Testing_Integration.ppt.

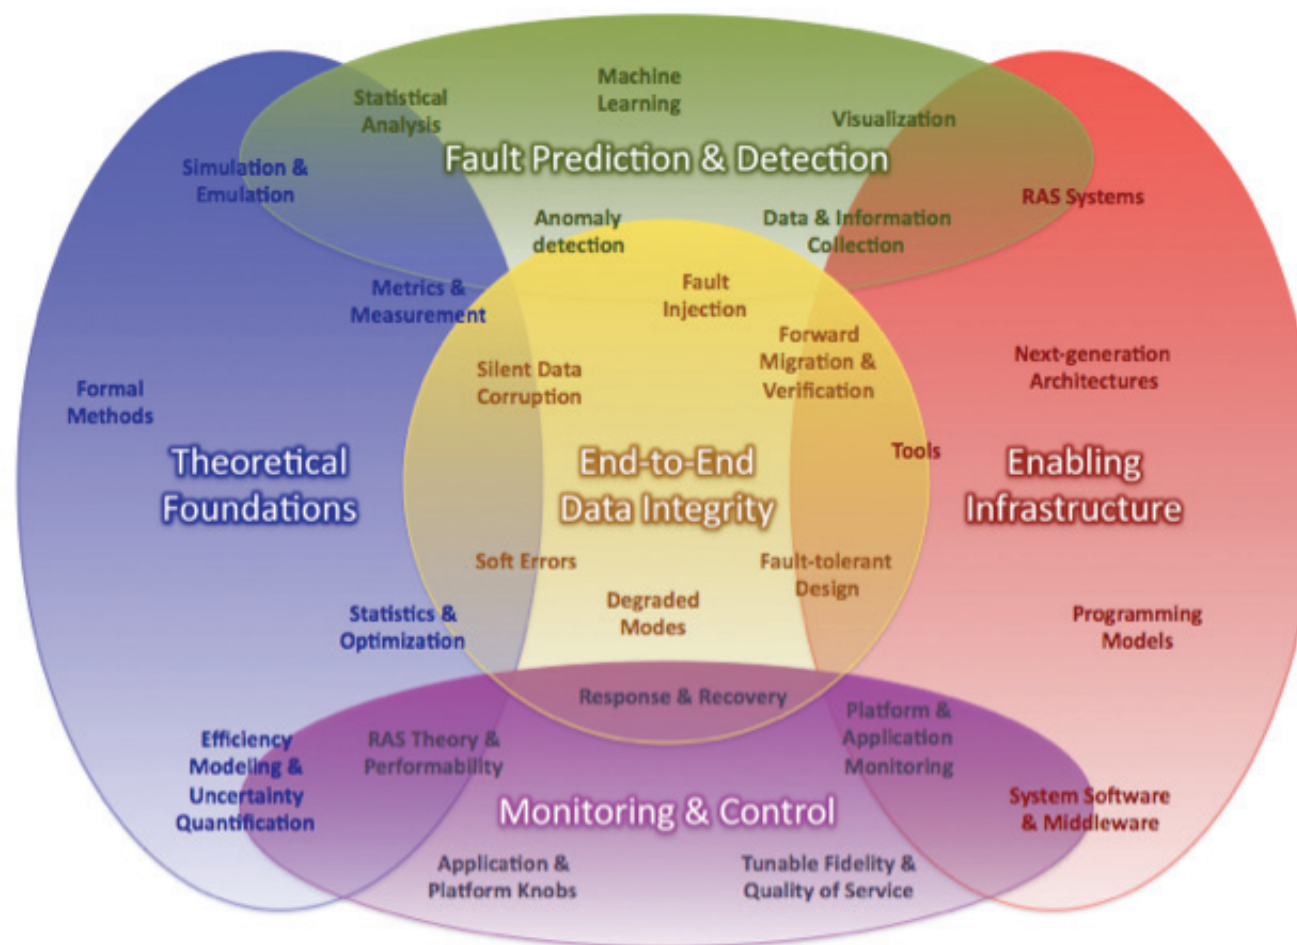


Fig. 1. Five focused but overlapping thrust areas.

For the full paper, see: http://institute.lanl.gov/resilience/docs/HECResilience_WhitePaper_Jan2010_final.pdf
 For links to the talks presented at the workshop see: <http://institute.lanl.gov/resilience/conferences/2009/>

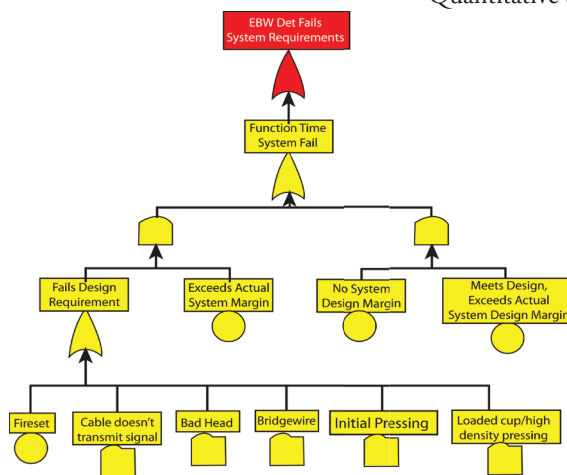
Statistical Models to Support Assessment and Decision-making in Stockpile Stewardship

Aparna V. Huzurbazar, Christine Anderson-Cook, Elizabeth Kelly, Michael Fugate, Todd Graves, Michael Hamada, Geralyn Hemphill, David Higdon, Richard Klamann, Lisa Moore, Richard Picard, Lawrence Ticknor, Joanne Wendelberger, Scott Vander Wiel, Brian Williams, CCS-6

The LANL Statistical Sciences Group (CCS-6) has a long-standing effort in research and applied methods to support the assessment of conventional and nuclear weapon stockpiles. Our work is focused on tracking and trending of stockpile parameter data streams, statistics-based quantification of margins and uncertainties (QMU), age-aware models for reliability with quantified uncertainty, and resource allocation. Our goal is to provide improved confidence in future weapons reliability, safety, and performance. Some of the projects that are currently supported are:

- Integrated Reliability Methodology Projects leveraged with the DOE/DOD Joint Munitions Program:
 - Munitions Stockpile Reliability Assessment
 - Complex System Health Assessment
- Surveillance Transformation
- Data Provenance
- Quantification of Margins and Uncertainties
- Quantitative Surveillance Metric Development

Fig. 1. Notional fault tree diagram for detonator performance.



We briefly describe selected work supported by these projects.

Integrated Reliability Methodology (IRM) supports a wide range of projects including system health assessment and software tools. One focus area of application has been developing methodology for new detonator designs. The approach uses fault trees (see Fig. 1) and reliability block diagrams to characterize detonator performance by the performance of key subcomponents [1]. This essentially breaks down component-level functionality

into the functionality, performance, and potential failure of events of key subcomponents. Having such information on components in existing and new designs allows us to better focus testing and helps determine priorities for improved understanding in the future.

Resource allocation methods are supported by IRM, and these methods mesh across elements in our work. Resource allocation focuses on assessing future data collection strategies for their projected improvement in the precision of system reliability estimates. This involves integrating the mathematical and computational methods necessary to connect integrated system assessments into a cost-benefit framework that can be used to compare different data collection strategies before the resources are spent and determine their anticipated utility. The results of this type of analysis can help determine how future resources should best be spent to understand system reliability and as part of improving weapon system surveillance planning (see Fig. 2)[2].

Surveillance transformation includes methodology for stockpile evaluation, age-aware assessments, and assessments of critical stockpile parameters. A major portion of the work is statistical modeling and evaluation for tracking and trending of surveillance data. In one specific application, statisticians are using advanced Bayesian hierarchical models for populations of shapes to model contours of the high-explosives parts in the stockpiles. This approach distinguishes between variation in as-built parts, instantaneous damage due to disassembly, and gradual aging (see Fig. 3). The approach also considers variation in the stockpile population, uncertainty in the stockpile averages, and measurement error. Analyses such as these can be used to provide input in the form of populations of contours under various assumptions, including extrapolating into the future, for the purpose of uncertainty quantification studies using physics codes [3].

Data provenance is concerned with managing large-scale, data-intensive projects that are a backbone of the weapons complex. The weapons complex collects a vast amount of data. Central to the ability to obtain and use the increasingly large amounts of data within the surveillance community is the need for tracking and

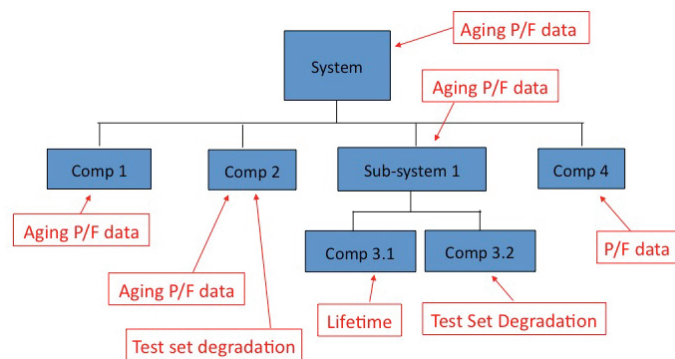


Fig. 2. Resource allocation for improved reliability with different types of data.

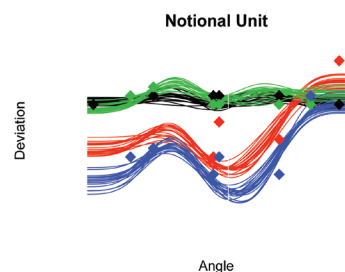


Fig. 3. Bayesian hierarchical modeling of population shapes for contours.

validation of the origins of the data. The term data provenance refers to the process of tracing and recording the origins of data and its movement between persons, locations, or databases. Focusing on the integrity of the data can allow for better evaluation of the value of the information and can help identify and/or resolve any discrepancies that may arise.

For surveillance data analysis,

it is essential to know exactly what data one is working with, where the data came from, whether or not the data have been validated, and any other information vital to the appropriate use of the data. This project worked with the Product Realization Integrated Digital Enterprise (PRIDE) program to develop a procedure to validate and track the evolution of data sets at LANL and at other DOE facilities. The team for the pilot project of this work received a Los Alamos Awards Program (LAAP) recognition. The project will facilitate the sharing of data within the nuclear weapons community.

Our QMU effort focuses on sensitivity and performance studies. We have developed methodologies for combining separate tests, and physics and engineering codes to map the effect of input variation and aging effects on performance [4]. We have created new methods for combining data from surveillance and experiments with physics simulations to provide performance-based lifetime assessments. The CCS-6 QMU work received a 2008 Defense Programs Award of Excellence.

We are involved with the development of a quantitative surveillance metric for confidence assessment of stockpile health [5]. The aging nuclear stockpile presents new challenges in our ability to develop and execute effective surveillance activities. Our work develops a tool to assist in the understanding of stockpile health. This is a statistical uncertainty model for reliability aging with an awareness of the consequences of sample size reduction and risks of various sampling rates for surveillance. This metric characterizes the uncertainty of stockpile reliability over time and can be used to help manage surveillance programs. We have also established methods for assessing the risks of sampling rates for surveillance [6] and quantifying reliability uncertainty, a joint project with Sandia National Laboratories (California and New Mexico) [7,8]. The goal is to integrate expert knowledge with summarized surveillance data. This approach supports balancing cost of data with precision of stockpile reliability estimates.

Through this broad range of applied work that encompasses developing methods and software tools, data collection strategies, and implementing solutions, CCS-6 supports

stockpile stewardship. The group works to apply best practices to improve the understanding of weapon assessment and to support decision-making about the enduring stockpile.

For more information contact
Aparna V. Huzurbazar at
aparna@lanl.gov.

- [1] D. Monroe et al., *An Element of Test Fire Optimization: Detonator Qualification Pilot Results (U)*, LA-CP-09-01225.
- [2] C. Anderson-Cook et al., *Qual. Reliab. Eng. Int.* **25**, 481 (2009).
- [3] T. Graves, *Trends and population variability in HE contours in the B61 (U)*, LA-CP 09-01326.
- [4] D. Higdon et al., *J. Am. Stat. Assoc.* **14**, 570 (2008).
- [5] S. Vander Wiel et al., *A Random Onset Model for Degradation of High Reliability Systems*, LA-UR 09-04667 (2009); also in *CASS Conf. Proc.* (NNSA in press).
- [6] M. Hamada et al., *Assessing the risks of sampling rates for surveillance*, LA-UR-09-06140.
- [7] A. Huzurbazar et al., *Reliability uncertainty aggregation case study for the B-61*, LA-UR 09-00172.
- [8] J. Lorio et al., *Quantifying Reliability Uncertainty: A Proof of Concept*, SAND-2009-2173.

Funding Acknowledgments

- DOE, NNSA, Enhanced Surveillance Campaign
- DOE, NNSA, Defense Programs and Core Surveillance
- DoD and DOE Joint Munitions Technology Development Program

Using the MIITS-Cyber Tool to Analyze Large-scale Cyber Threats

Guanhua Yan, Stephan Eidenbenz, CCS-3

As computer networks have permeated almost every aspect of the nation's critical infrastructure such as its transportation, power grid, communications, and defense, the importance of ensuring a secure and robust information technology (IT) infrastructure is tremendous. The direct costs of cybercrime to the US economy reach tens of billions of dollars annually [1], and the federal government spends billions of dollars on cybersecurity each year [2]. Given the increasing intensity of cyber threats, the Obama administration has recently claimed cyber security a national security priority.

A virtual controllable testbed is an indispensable tool to evaluate the effect of potential cyber threats and the effectiveness of proposed countermeasures. As large-scale cyber attacks (malware, botnet, and distributed denial of service [DDoS] attacks) commonly

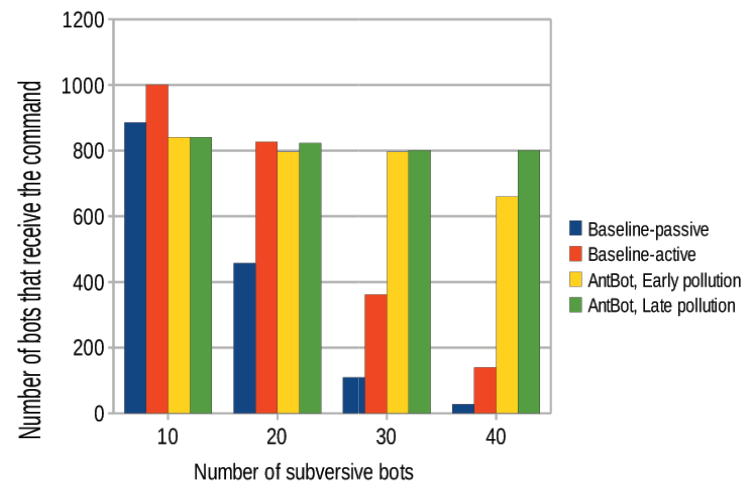
involve many machines, such an evaluation testbed should not only capture the realism necessary for the analysis, but also have the merit of high scalability. Towards this end, we extended the capabilities of the LANL-developed modeling and simulation suite MultiScale Integrated Information and Telecommunications System (MIITS) and studied two emergent large-scale cyber threats: botnets, which use the peer-to-peer (P2P) networks

for their command and control (C&C), and malware, which spread in large cyber social networks (e.g., Twitter and Facebook). The simulation modules that fulfill the analysis of these two cyber threats are called BotSim and CyberSim, respectively.

BotSim. To achieve great realism in modeling P2P behaviors by P2P-based botnets, we use the actual development code of a popular P2P client, aMule, which is based on the KAD protocol, a variant of Kademlia [3]. It is known that P2P-based botnets, although they do not suffer a single point of failure as IRC-based botnets do, can be effectively disrupted by polluting the command keys [4]. We, however, predict potential moves by the botmaster and find that a new type of hypothetical botnet (dubbed AntBot) works resiliently against pollution-based mitigation [5]. The key idea is to use a tree-like multilevel structure to relay C&C messages from the botmaster in P2P networks. Figure 1 shows the resilience of this hypothetical P2P-based botnet against pollution-based mitigation for a 1000-node botnet. The baseline cases refer to the botnets that do not deploy the antipollution scheme. From the figure, it is clear that even with a number of subversive bots controlled by the white-hat defender, the majority of the bots can still obtain the command issued by the botmaster, regardless of the pollution scheme used by the white-hat defender. The take-home message from this study is that cyber defenders should not be content with their capability of defending against past observed cyber attacks, but should also deploy proactive defense systems that prevent future unobserved cyber attacks from happening.

CyberSim. Virus propagation in social networks has been intensely studied in the literature, especially from a structural perspective. For online social networks, virus propagation is affected by user activities, such as when the user goes online. From a real-world trace, we find that the number of activity events generated by each user in an online social network is well characterized by an extended exponential distribution. The highly skewed distribution of online user activities significantly affects how virus propagates in online social networks. We developed CyberSim, which can be driven by activity traces of real-world online social network users, to quantify

Fig. 1. Resilience of AntBot against pollution-based mitigation.



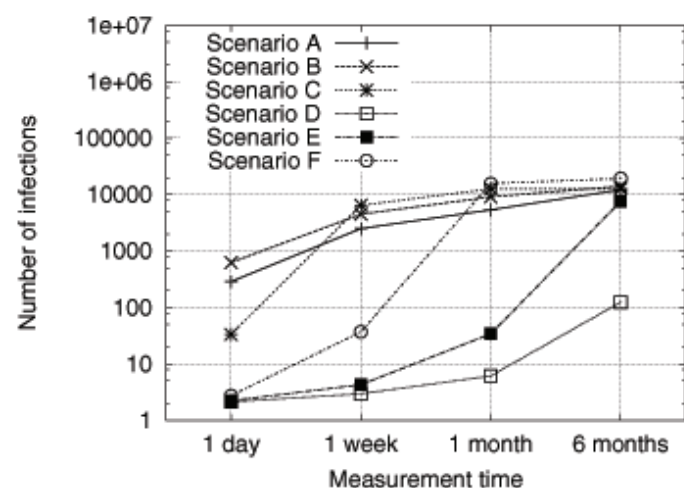


Fig. 2. Virus propagation under different scenarios.

the impact of social structures and user activity patterns on virus propagation [6]. Using CyberSim, we investigate virus propagation under six different scenarios—Scenarios A, B, and C use a real-world social structure, and Scenarios D, E, and F use the Erdos-Renyi social structure; Scenarios A and D use real-world activity events; Scenarios B and E randomize the time intervals among activity events by each user; Scenarios C and F randomize the time intervals among activity events across all the users. Figure 2 depicts the propagation progress by the virus under different scenarios.

From the results, we conclude that the realistic social structure, which has been shown to be a power-law and small-world graph, spreads the virus quickly at its early stage compared with the Erdos-Renyi social structure. Also, we see that the highly skewed distribution of activity events among online users in the real-world dataset actually slows down the virus propagation process. This suggests that future work on analyzing virus propagation in social networks should consider not only realistic social graphs but also realistic human activity models.

For more information contact Guanhua Yan at ghyang@lanl.gov.

- [1] 2005 FBI/CSI computer crime survey, <http://www.gocsi.com>, January (2006).
- [2] Pulliam, D. "Cybersecurity spending estimated to grow to \$7.1 billion by 2009," *Government Executive*, <http://www.govexec.com/dailyfed/0305/031705p1.htm> (2005).
- [3] D.T. Ha et al., "On the Effectiveness of Structural Detection and Defense against P2P-based Botnets," *Proc. 39th Ann. IEEE/IFIP Int. Conf. Dependable Systems Networks (DSN'09)* (2009).
- [4] T. Holz et al., "Measurements and Mitigation of Peer-to-Peer-based Botnets: A Case Study on Storm worm," *Proc. 1st Usenix Workshop Large-Scale Exploits and Emergent Threats (LEET'08)* (2008).
- [5] G. Yan, D.T. Ha, S. Eidenbenz, *AntBot: Anti-Pollution Peer-to-Peer Botnets*, LA-UR 09-06004.
- [6] G. Yan et al., *Towards A Deep Understanding of Malware Propagation in Online Social Networks*, LA-UR 09-08100.

Funding Acknowledgments

DHS, National Infrastructure Simulation and Analysis Center (NISAC) Program

Materials Science

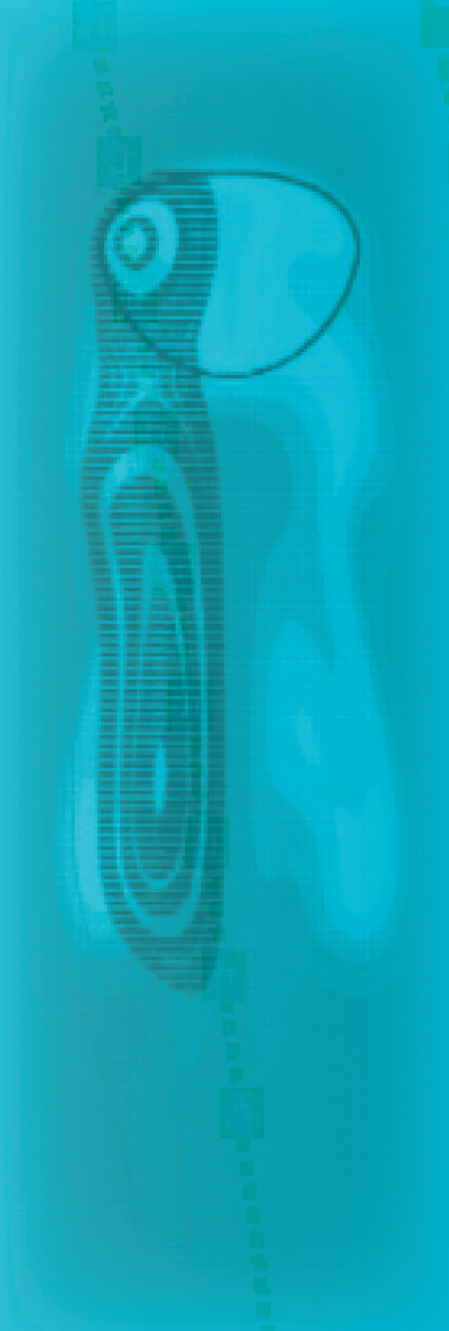
This section is representative of the breadth and depth of work in materials science in the Theory, Simulation, and Computing Directorate. The work represented here includes applications to organic electronic materials, mechanical properties of polycrystalline metals, materials under extreme conditions, fuel reprocessing, aging and lifetime prediction, the response to ballistic projectiles, the mechanical behavior of nanowires, and the properties of nuclear fuels.

The techniques employed span a broad range both in the physical modeling approaches and in the mathematical and numerical algorithms used in the simulations. At atomic and mesoscale dimensions we see the use of quantum mechanical and molecular dynamics. This includes one example that develops theoretical methods to combine the two. At the microscale and continuum end we learn about traditional methods from solid

mechanics, including dislocation dynamics, shock wave techniques, fluid dynamics, heat transport, as well as statistical analyses.

The combination of such broad capabilities with the talent of our personnel make ADTSC a world-class organization in materials modeling and simulation.

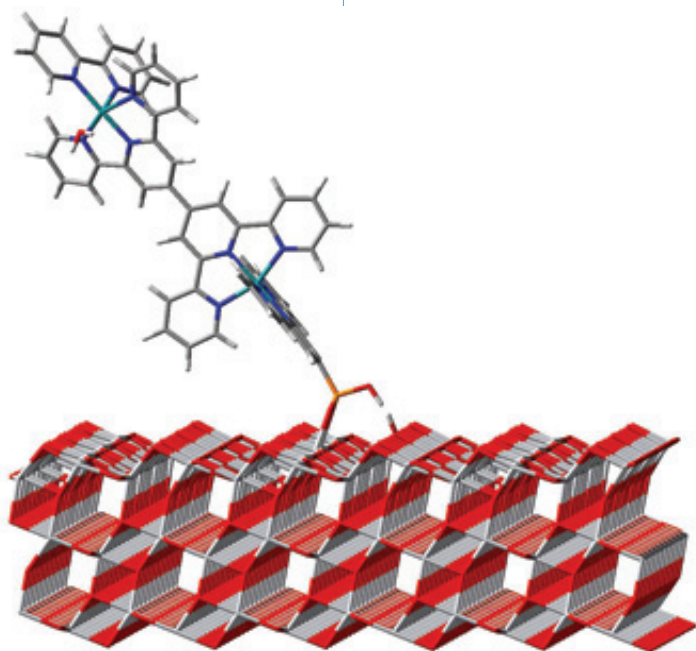
Materials Science



Systematic Study of Modifications to Ru(II)-polypyridine Dyads for Electron Injection Enhancement

Elena Jakubikova, Richard L. Martin, Enrique R. Batista, T-1; Robert C. Snoeberger III, Victor S. Batista, Yale University

Fig. 1. $[(bpy)(H_2O)Ru(tpy-tpy)Ru(tpy)]^{4+}$ attached to (101) surface of anatase TiO_2



Transition metal complexes, and particularly ruthenium polypyridine complexes, can be used as catalysts or photocatalysts capable of performing water oxidation [1] or oxidation of organic compounds.[2,3] They can also be used as chromophores to harvest solar energy due to their absorption in the visible region. Therefore, understanding the nature of the electronic excited states in transition-metal complexes is crucially important for elucidation of the mechanistic details of the photocatalyst function and design of efficient molecular devices for solar energy conversion.

An important aspect of the molecular photocatalyst functionality is its ability to inject electrons into a metal oxide nanoparticle in order to induce charge separation and thus mimic photosynthetic charge transfer events [4-6]. This effect is usually achieved by an excitation into a metal-to-ligand charge transfer (MLCT) state that couples to the states in the nanoparticle conduction band, and this coupling drives the interfacial electron transfer (IET). IET competes with other processes that occur upon photoexcitation, such as radiative or nonradiative

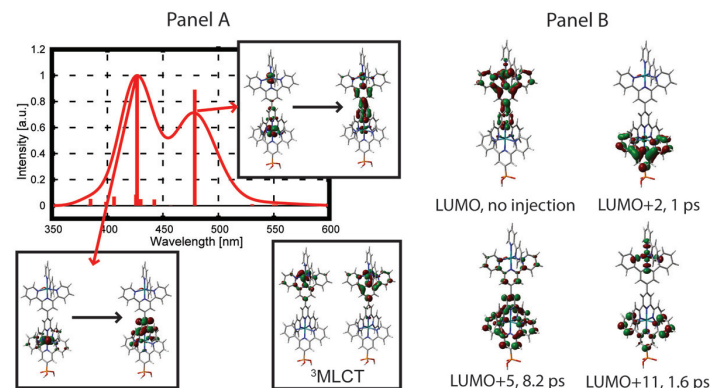


Fig. 2. Panel A: Calculated absorption spectrum of $[(bpy)(H_2O)Ru(tpy-tpy)Ru(tpy-PO_3H_2))]^{4+}$ with natural transition orbitals for the most intense transitions and singly occupied natural orbitals of the optimized 3MLCT state for the same molecule. Panel B: Selected virtual molecular orbitals of $[Ru(tpy)(bpy)(H_2O)-Ru(tpy)(tpy(PO_3H_2))]^{4+}$ adsorbate obtained from extended Hückel theory, and their IET rates.

transition back to the ground state, or intersystem crossing (IC) into the lowest triplet excited state. IC plays an important role in the case of the ruthenium polypyridine complexes [7], and the excited states that lead to the electron injection into the semiconductor are a combination of initially populated singlet excited states and a thermalized 3MLCT state [8].

We have studied absorption properties and IET dynamics in a prototype photocatalytic assembly [3] $[(bpy)(H_2O)Ru(tpy-tpy)Ru(tpy)]^{4+}$ ($[Ru(tpy)(bpy)(H_2O)]^{2+}$ - model catalyst, $[Ru(tpy)_2]^{2+}$ - chromophore; tpy = 2,2':6',2''-terpyridine and bpy = 2,2'-bipyridine) attached to an anatase TiO_2 nanoparticle via a phosphonic acid linker. Density functional theory (DFT) was used to obtain the ground-state geometry of the catalyst-chromophore-nanoparticle assembly (see Fig. 1), as well as the absorption spectrum and the lowest triplet excited state of $[(bpy)(H_2O)Ru(tpy-tpy)Ru(tpy-PO_3H_2))]^{4+}$. Wave packet quantum dynamics simulations based on extended Hückel Hamiltonian [9] were used to obtain the IET rates from the excited states localized on the catalyst-chromophore assembly into the nanoparticle.

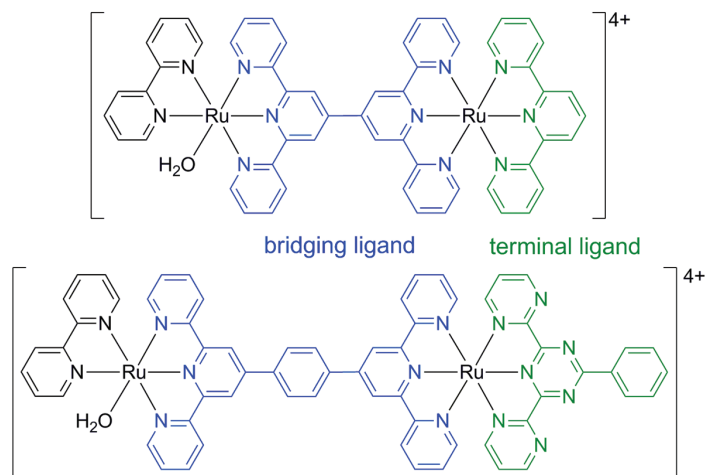


Fig. 3. Chemical structure of the original catalyst-chromophore dyad (top) and the modified catalyst-chromophore dyad (bottom).

The absorption spectrum obtained using the time-dependent DFT formalism in the visible region for the free $[(bpy)(H_2O)Ru(tpy-tpy)Ru(tpy-PO_3H_2)]^{4+}$ molecule is shown in panel A of Fig. 2. The two most intense peaks correspond to the excitation of the electron into the orbitals characterized by the electron density on the bridging tpy-tpy ligand. Natural transition orbitals, which describe these excitations, are also shown in Fig. 2. The excited $[(bpy)(H_2O)Ru(tpy-tpy)Ru(tpy-PO_3H_2)]^{4+}$ molecule can also undergo intersystem crossing into the 3MLCT state, in which the excited electron localizes on the tpy-tpy bridging ligand as well (also shown in panel A of Fig. 2). Panel B of Fig. 2 shows selected virtual molecular orbitals of the $[(bpy)(H_2O)Ru(tpy-tpy)Ru(tpy-PO_3H_2)]^{4+}$ adsorbate obtained with the extended Hückel theory, and their IET rates.

While the IET occurs at a 1-10 ps rate from the orbitals localized on the terminal tpy(PO_3H_2) ligand attached to TiO_2 , the most intense transitions, as well as long-lived 3MLCT state, are characterized by the spatial localization of the excited electron on the bridging tpy-tpy ligand. The bridging tpy-tpy ligand is only weakly coupled with the TiO_2 nanoparticle, which leads to an inefficient IET in the investigated catalyst-chromophore-nanoparticle assembly. Therefore,

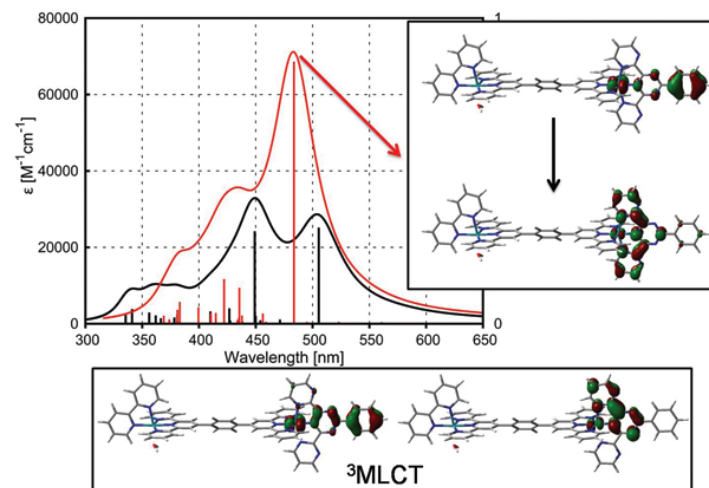


Fig. 4. Top panel: Absorption spectra of the original dyad (black) and modified dyad (red) and the natural transition orbitals corresponding to the most intense excitation in the visible region. Bottom panel: Lowest triplet MLCT state of the modified dyad.

we explored a series of 18 modifications to the bridging and terminal ligands, based on previously synthesized molecules, in order to find a related catalyst-chromophore assembly that would favor electronic excitations into the terminal terpyridine ligand with a strong electronic coupling to the semiconductor [10].

After investigating various modifications, we found that both introduction of a spacer group (such as phenylene or alkane) into the tpy-tpy bridge and replacement of the terminal terpyridine group by a more extended π -conjugated ligand are necessary to capture the excitations next to the attachment group that links to the surface. An example of a successfully modified dyad is shown in Fig. 3. The modified dyad displays both intense absorption into the terminal heteropyridine ligand, and electron localization on the

modified terminal ligand in the lowest 3MLCT state, which should lead to the efficient IET upon attachment to the TiO_2 nanoparticle.

In conclusion, we have shown that the excited state properties of the prototype catalyst-chromophore dyad can be tuned by the use of different bridging ligands and modifications to the terminal ligand. These results have implications for the design of photocatalysts and dye-sensitizer assemblies based on Ru(II)-terpyridine compounds.

**For more information contact
Enrique R. Batista at
erb@lanl.gov.**

- [1] J.K. Hurst, *Coord. Chem. Rev.* **249**, 313 (2005).
- [2] J.A. Treadway, J.A. Moss, T.J. Meyer, *Inorg. Chem.* **38**, 4386 (1999).
- [3] W. Chen, et al., *Chem., Int. Ed.* **48**, 9672 (2009).
- [4] L.C. Sun et al., *Chem. Soc. Rev.* **30**, 36 (2001).
- [5] M. Falkenstrom, O. Johansson, L. Hammarstrom, *Inorg. Chim. Acta* **360**, 741 (2007).
- [6] R. Ghanem et al., *Inorg. Chem.* **41**, 6258 (2002).
- [7] J.P. Sauvage *Chem. Rev.* **94**, 993 (1994).
- [8] C. She et al., *J. Phys. Chem. A* **111**, 6832 (2007).
- [9] L.G.C. Rego, V.S. Batista, *J. Am. Chem. Soc.* **125**, 7989 (2003).
- [10] E. Jakubikova, R.L. Martin, E.R. Batista, *Inorg. Chem.* **49**, 2975 (2010).

Funding Acknowledgments

LANL Directed Research and Development Program

Modeling and Characterization of Grain Scale Strain Distribution in Polycrystalline Tantalum

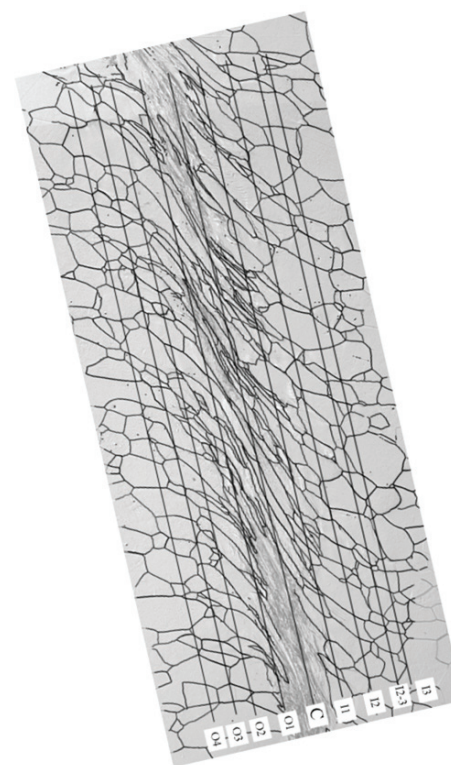
Curt A. Bronkhorst, T-3; Amy R. Ross, MST-16; Benjamin L. Hansen, Hashem M. Mourad, T-3; Ellen K. Cerreta, John F. Bingert, MST-8

A common sample geometry used to study shear localization is the tophat, an axi-symmetric sample with an upper hat portion and a lower brim portion. The gage section lies between the hat and brim. The gage section length is on the order of 0.9 mm with deformation imposed through a Split-Hopkinson Pressure Bar system at maximum top-to-bottom velocity in the range of 10-25 m/s. Detailed metallographic analysis has been performed on sections of the samples to quantify the topology and deformation state of the material after large deformation shear, shown in Fig. 1. These experiments, performed with polycrystalline tantalum, have been modeled using a multiscale, polycrystal plasticity approach. A Voronoi tessellation-based microstructural model and a coupled thermo-mechanical elasto-viscoplastic crystal plasticity model were used. The crystal plasticity model allowed for slip to occur on the twelve $\{110\}\langle 111 \rangle$ and twelve $\{112\}\langle 111 \rangle$ slip systems. Three numerical models were produced using three different realizations of initial crystallographic texture distribution within the same morphological microstructure and the results presented. The results of one of these realization simulations are shown here. The detailed metallographic analysis of the deformed sample shear zone produced an estimate for the strain profile within that region, and these results are compared directly with the three numerical simulation results, given in Fig. 3. The experiments produce lower and upper bound estimates of the deformed strain state in the material due to ambiguity of the initial state for the 2D metallographic technique available. Although the stress response of the models predict a stress response that is greater than that observed experimentally, the local strain response compares very well with the results of the metallographic analysis. Both

the experimental and numerical results give estimates of the degree of heterogeneity that exists in the deformation response of metallic polycrystalline aggregates. This heterogeneity is believed to drive the initiation of damage and failure processes.

For more information contact Curt A. Bronkhorst at cabronk@lanl.gov.

Fig. 1. Digitized shear zone region highlighting the grain boundary structure and used to quantify the deformation profile. Note that measurements were not taken along line I2-3.



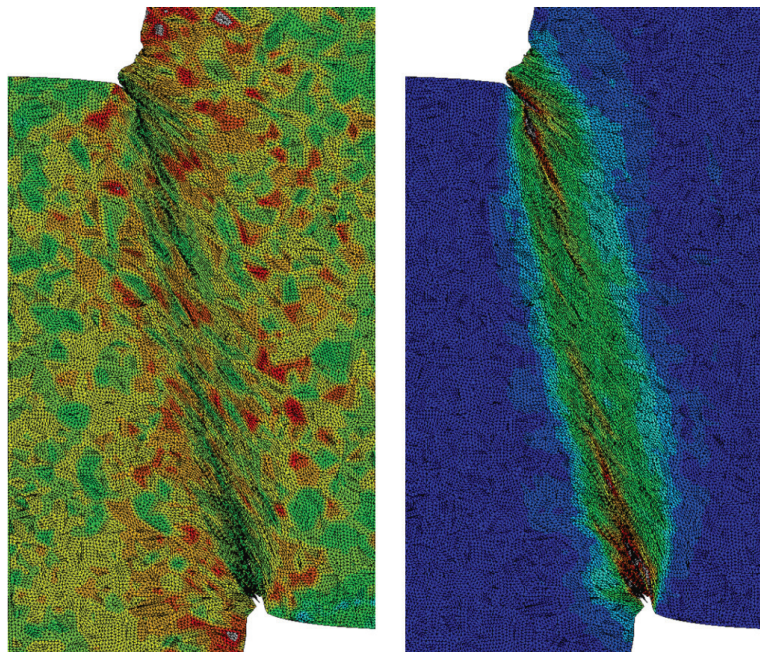


Fig. 2. vonMises stress (left, red=900 MPa) and equivalent plastic strain (right, red=1.5) contours in the deformed mesh of the realization 1 model.

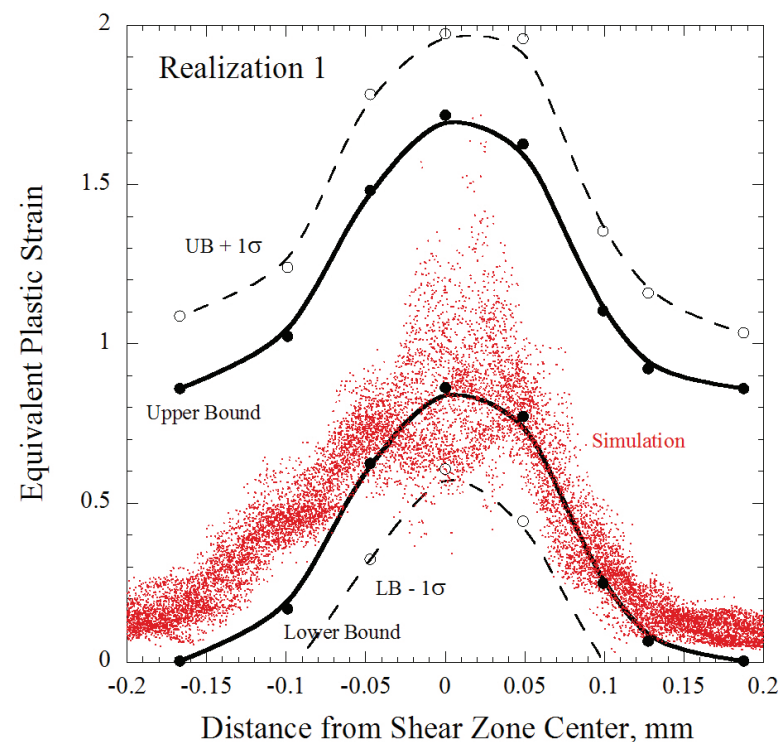


Fig. 3. Comparison of the results of realization 1 simulation to the strain evaluated from grain aspect ratio measurements. Negative position values indicate points nearest to the center of the axis-symmetric sample. The solid curves represent the mean lower and upper bound equivalent plastic strain evaluated from the metallographic images. The dashed curves represent the one standard deviation envelope around the experimental lower and upper bound curves. The red points represent the numerical results from all material points within the extracted shear zone region.

Funding Acknowledgments

DOE, NNSA Advanced Simulation and Computing (ASC) Program

Nonequilibrium Volumetric Response of Shocked Polymers

Bradford E. Clements, T-1

Consider a high molecular weight polymer. Each polymer molecule may contain hundreds to even millions of atoms. The backbone of the molecular structure is the polymer chain. For amorphous polymers these chains typically form complex networks. It is common for the polymer to have physical entanglements, chemical cross linkages, or both. Because of the inherent complexity at the molecular level, it is relatively easy to drive an equilibrated polymer into a nonequilibrium state by application of a simple mechanical deformation. The externally applied deformation can be shear, or more generally deviatoric, but also volumetric in nature.

It is widely recognized that the chain mobility is the physical quantity that governs how successful a polymer will be at relaxing back into an equilibrium state for the deformed system. The chain mobility is the ability of the polymer molecules to locally rearrange themselves, to bring the deformed structure back into a state of equilibrium. Obviously the available thermal energy is an important factor in this process. For temperatures far above the glass transition temperature T_g , the chain mobility is high and a polymer is expected to relax quickly back into a state of equilibrium. However, below T_g the chain mobility is low and the polymer system may effectively become frozen into a long-lived metastable nonequilibrium state.

In any mechanical experiment, besides the loading rate, two other characteristic times are important when considering polymer relaxation—the total duration of the experiment and the experimental resolution over which the measurements are taken. Thus, first consider a simple low rate uniaxial compression experiment used to measure the polymer's stress-strain behavior. For discussion it is useful to make the gross simplification that molecular relaxation processes can be categorized as fast, intermediate, and slow. This categorization corresponds approximately to a

prototypical shear relaxation spectrum characteristic of many amorphous polymers [1]. Fast processes involve only few atoms (or chain segments) in single molecules and are associated with high frequency vibrational modes. At the other extreme are the relaxation modes associated with the collective motion of many chain segments involving many neighboring polymer molecules. For an entangled polymer, relaxation times for these processes are correspondingly very long. Between these extremes are the numerous intermediate rate relaxations. When a uniaxial compression experiment is done the system is immediately driven out of equilibrium. Fast relaxations occur so rapidly that such modes fully relax long before the first measurement is taken. However, long-lived modes will continue over the duration of the experiment.

The process of polymer relaxation and its effects on the mechanical behavior is encompassed in the theory of polymeric viscoelasticity, and because the viscoelastic response is associated mainly with the deviatoric response, the relevant theory is deviatoric viscoelasticity. When the range of strain rates is limited, as in the case of laboratory uniaxial compression experiments, volumetric viscoelasticity may typically be ignored. One then invokes an equilibrium equation of state (EOS) to describe the volumetric behavior. A more fundamental reason for ignoring volumetric viscoelasticity is that for small volumetric changes, it is primarily the high frequency vibrational modes (those associated with fast relaxation processes) that are likely to be compatible with the volumetric deformation, and thus likely to be excited. As long as the experiment is long and the resolution is low, relaxation of these modes will not be observable.

Next, consider a plate impact experiment where, by the action of a shock introduced into the system, the polymer will rapidly transition from an initial equilibrium state to a high-pressure shocked state. The duration of these experiments is typically microseconds (before release waves interrupt the shocked state), and the resolution is on the order of nanoseconds. For this case we argue that volumetric relaxations from the nonequilibrium shocked state to the equilibrium shocked state, i.e., volumetric viscoelasticity, must not be ignored.

To appreciate the deviation of the shocked state from the equilibrium state, we first formulate the equilibrium EOS. Here we will define the equilibrium EOS as that which is obtained using slow measurements, i.e., those lasting minutes to hours. Typical measurements satisfying this criterion are the calorimetric-measured specific heat and dilatometry-measured specific volume. Both of these are available in the literature—a tabulated specific heat for polymethyl-methacrylate (PMMA) is listed in the Advanced Thermal Analysis System (ATHAS) database [2], and specific volume isobars are reported in [3]. To use this data a semi-empirical expression for the Gibbs free energy is parameterized. The resulting Gibbs free energy (Fig. 1), and isothermal bulk modulus (Fig. 2) are shown for illustration.

Because the Gibbs free energy describes a complete EOS, the equilibrium Hugoniot can be calculated by satisfying the Rankine-Hugoniot jump conditions. The resulting equilibrium Hugoniot is plotted using the shock and particle velocities and compared with the shock experimental Hugoniot of [4]. It is immediately obvious from Fig. 3 that the two Hugoniots show major discrepancies. We speculate that this difference is due to the volumetric response of a polymer being rate dependent; while the equilibrium EOS measurements have deformation rates occurring over inverse minutes to hours, the shock measurements have a characteristic deformation rate of 10^6 - 10^7 s⁻¹. We thus assert that the bulk relaxation function should increase by about a factor of two as PMMA is shocked from equilibrium up to shock loading rates, i.e., volumetric viscoelasticity must not be ignored.

To include nonequilibrium effects, a Helmholtz free energy is written as the sum of an equilibrium, A_e , plus a nonequilibrium contribution coming from the deviatoric "D" and bulk "B" parts [5]:

$$\begin{aligned} \rho A(\varepsilon_v, e_{ij}^e, T, \dot{\varepsilon}_v, \dot{e}_{ij}^e) &\approx \rho A_e(\varepsilon_v, e_{ij}^e, T) \\ &+ \int_0^t \int_0^t G_D(t_1, t_2; t) \frac{de_{ij}^e(t_1)}{dt_1} \frac{de_{ij}^e(t_2)}{dt_2} dt_1 dt_2 \\ &+ \frac{1}{2} \int_0^t \int_0^t G_B(t_1, t_2; t) \frac{d\varepsilon_v(t_1)}{dt_1} \frac{d\varepsilon_v(t_2)}{dt_2} dt_1 dt_2 \end{aligned}$$

Here, $\varepsilon_v, e_{ij}^e, T, \dot{\varepsilon}_v, \dot{e}_{ij}^e, G_D$ and G_B are the volumetric strain, deviatoric elastic strain, temperature, volumetric strain rate, deviatoric strain rate, deviatoric kernel, and volumetric kernel, respectively. Determining functional expressions for G_D and G_B , and evaluating the resulting pressure, deviatoric stress, dissipation, and entropy is a long and rather complex process, the details of which are unimportant here. Finally, the model is implemented in a finite element code to perform simulations.

In the plate impact experiments of [4], 6.35mm PMMA plates impacted 6.35mm PMMA plates at various impact speeds. Figure 4 shows the experiment and our theoretical calculated shock profiles. Clearly the comparison is very good. Similarly, the theoretical shock Hugoniot now also agrees with that of Barker and Hollenbach [4]. The theory also shows that PMMA never reaches true equilibrium during the shock experiment.

For more information contact Bradford E. Clements at bclements@lanl.gov.

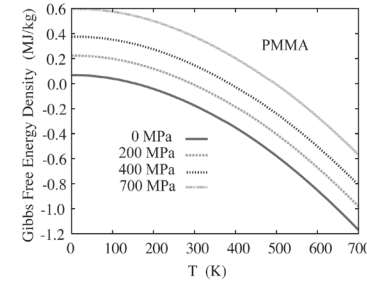


Fig. 1. PMMA equilibrium Gibbs free energy.

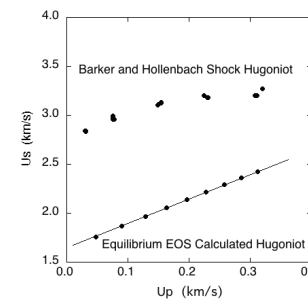


Fig. 3. Shock Hugoniot from [2] and that calculated from the equilibrium EOS.

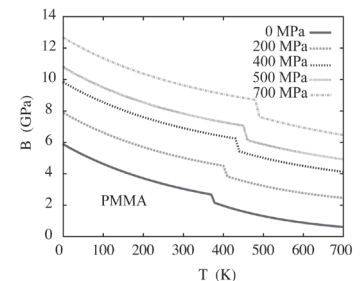


Fig. 2. PMMA equilibrium isothermal bulk modulus.

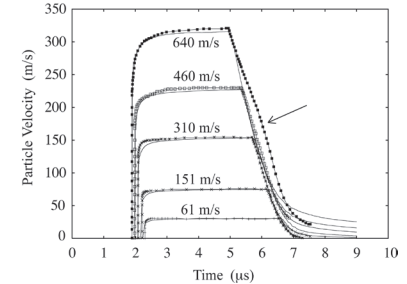


Fig. 4. PMMA shock profiles [2] (points) and from the theory (lines). Impact velocities are listed.

- [1] J.D. Ferry, *Viscoelastic Properties of Polymers*, Wiley, New Jersey (1970).
- [2] ATHAS (The Advanced Thermal Analysis System), <http://athas.prz.rzeszow.pl/>
- [3] S. Theobald, W. Pechhold, B. Stoll, *Polymer* **42**, 289 (2001).
- [4] L.M. Barker, R.E. Hollenbach, *J. App. Phys.* **41**, 4208 (1970).
- [5] R. M. Christensen, *Theory of Viscoelasticity, An Introduction*, Academic Press (1971).

Funding Acknowledgments
DOE/DoD Joint Munitions Program

Material Dynamics at Extreme Conditions

Timothy C. Germann, T-1; Sheng-Nian Luo, P-24; Sriram Swaminarayan, CCS-2

Fig. 1. Incipient spall failure in a copper bicrystal (shocked left-to-right) by homogeneous void nucleation along the (vertical) plane of maximum tensile stress. Only undercoordinated surface atoms are shown, so that one sees the free surfaces of the impactor and target and the voids that have begun to nucleate, grow, and coalesce. The (horizontal) grain boundaries, one in the center of the sample and the other at the top/bottom periodic boundary, remain intact and are not visible.



Dynamic loading, such as high-speed sliding friction or shock impact, can dramatically affect the microstructure and properties of materials on ultrafast timescales that are difficult—if not impossible—to probe experimentally. Because sound waves travel at a few km/s, or equivalently a few nm/ps, through a typical metal, they traverse the sub-nm interatomic lattice spacing in much less than a picosecond (ps). While experimentally challenging, processes occurring at such nm length and ps time scales are ideal for study by nonequilibrium molecular dynamics (MD) simulations. Over the past decade, large-scale MD simulations have provided significant insight into the microscopic pathways and kinetics of shock-induced plasticity and phase transformations in single-crystal metals. Sample

sizes of a few million atoms are typically sufficient; representing a cube of material with edge lengths 10-100 nm, this can capture the emergent length scale (average spacing between dislocations or product phase nuclei). Such simulations first suggested that the polymorphic (bcc-hcp) transformation in shocked iron can take place on ps timescales [1]. The predicted orientation relationship, timescale, and product grain size were all subsequently confirmed by ultrafast in situ X-ray diffraction measurements on laser-shocked iron thin foils [2].

The SPaSM (Scalable Parallel Short-range Molecular Dynamics) code was originally developed in the early 1990s for the newly emerging era of massively parallel supercomputers such as the Thinking Machines CM-5, and achieved IEEE Gordon Bell Prize-winning performance by minimizing memory usage and floating-point operations. However, for the modern generation of heterogeneous, multicore architectures such as Roadrunner, it is data movement rather than storage or arithmetic operations that is increasingly the bottleneck. Thus, we have redesigned the entire communication infrastructure and data structures of SPaSM to allow for better asynchrony between the processors, in particular to accommodate the IBM Cell BroadBand Engine (Cell BE) processor and the unique multilayer hierarchy of Roadrunner's architecture. This effort has paid off, resulting in double-precision benchmark performance of 369 TFlop/s on the full machine [3], and a ~5x speedup for the embedded atom method (EAM) potentials typically used to model simple metals such as copper, silver, and iron.

Using the SPaSM code on Roadrunner, we are investigating the ejection of material that can occur from shocked surfaces [4]. The goal of this work is to develop models that can predict the amount of mass ejected from a shocked interface with a given surface finish and loading history (peak shock pressure, either from a supported square-wave or explosive-driven Taylor wave). We would also like to understand how that mass is distributed, namely its particle size and velocity distributions, as well as evolutions and correlations, if any, between the two. Experimentally, the total mass can be inferred by measuring the resulting momentum transfer onto an Asay foil or piezoelectric probe at some standoff distance, while particle sizes larger than a micron can be imaged using holography or X-ray radiography. However, there is no direct experimental information on the distribution of particle sizes smaller than a micron, nor on the correlation between size and velocity distributions.

On such microscopic scales, MD simulations can complement experiments by providing unique insight into the material dynamics at submicron length and subnanosecond timescales, including key fragmentation and atomization mechanisms, but until now this problem had remained computationally intractable. The

formation and transport of ejecta involves a complex range of physical processes, including Richtmyer-Meshkov instability (RMI) development in solid materials with dynamic material strength properties, classical and turbulent fragmentation and atomization, and particulate transport in a turbulent gas. We have carried out a systematic study of RMI development from a single sinusoidal surface perturbation in copper to test various RMI theories including material strength effects, which suppress the instability growth. We are using these simulations to study the evolution of the density and velocity distributions of the ejected mass, the modes of particle breakup, and ultimately to develop source theories of ejecta formation based on RMI growth, including material strength effects and transport models that describe the temporally evolving particle size and velocity distributions. Earlier MD simulations were able to demonstrate the initial jet formation but could not reach timescales long enough to observe the subsequent necking instabilities leading to jet breakup and droplet formation that have now been revealed (Fig. 1). These fragmentation and atomization processes are also difficult to study experimentally, although various theories have been proposed; atomistic-level simulations such as those presented here are contributing to the development of physics-based models at LANL.

Ejecta is only one form of shock-induced material failure, occurring when a shock wave reflects from a free surface to become an expanding rarefaction (or release) fan. When two such rarefaction fans (one from the impactor free surface, the other from the target) intersect in the interior of the material, they put the material into tension and can lead to spall failure [5]. Ductile spall failure results from the nucleation, growth, and coalescence of voids; models have been developed that account for each of these aspects, but without direct experimental information; due again to the ultrafast time and ultrasmall length scales. Using Roadrunner, we have been able to study this process in copper bicrystals, revealing the competition between heterogeneous void nucleation at defects such as grain boundaries (GB) and homogeneous nucleation within the bulk single crystal. Large system sizes are required to clearly separate the two processes, and long timescales to explore a wide range of

strain rates. Figure 2 shows a Cu bicrystal loaded parallel to the horizontal GBs (one in the center and the other at the top/bottom periodic boundary), with 400 m/s impact velocity. The sample is 230 nm tall (i.e., each grain is 115 nm), 20-nm thick (into the plane, i.e., images are looking through the entire 20-nm thickness), and 205-nm long (54 million atoms in all). Following shock compression and release, dislocations and voids are produced that leave the sample in an incipient spall state; that is, with a number of voids that remain intact after growing and coalescing, but that have not caused complete fragmentation. In this example, the short sample length leads to a very high strain rate, with voids primarily nucleated homogeneously (within the grains) along the vertical plane of maximum tensile stress. On the other hand, a longer sample length (1 μm , 270 million atoms) leads to a lower strain rate and sufficient time for void nucleation to be localized at the GBs, changing the failure mode from a vertical spall plane to a horizontal grain decohesion. These results indicate the interplay between grain size and the timescale for nucleation kinetics, with a competition between heterogeneous and homogeneous nucleation.



Fig. 2. Incipient spall failure in a copper bicrystal by heterogeneous void nucleation along the (horizontal) grain boundaries.

For more information contact
Timothy C. Germann at
tcg@lanl.gov.

- [1] K. Kadau et al., *Science* **296**, 1681 (2002).
- [2] D.H. Kalantar et al, *Phys. Rev. Lett.* **95**, 075502 (2005).
- [3] T.C. Germann, K. Kadau, S. Swaminarayan, *Concurrency Computat.: Pract. Exper.* **21**, 2143 (2009).
- [4] T.C. Germann et al., "Large-scale molecular dynamics simulations of particulate ejection and Richtmyer-Meshkov instability development in shocked copper," in *DYMAT 2009: 9th International Conference on the Mechanical and Physical Behaviour of Materials under Dynamic Loading*, 1499 (2009).
- [5] S.-N. Luo, T.C. Germann, D.L. Tonks, *J. Appl. Phys.* **106**, 123518 (2009).

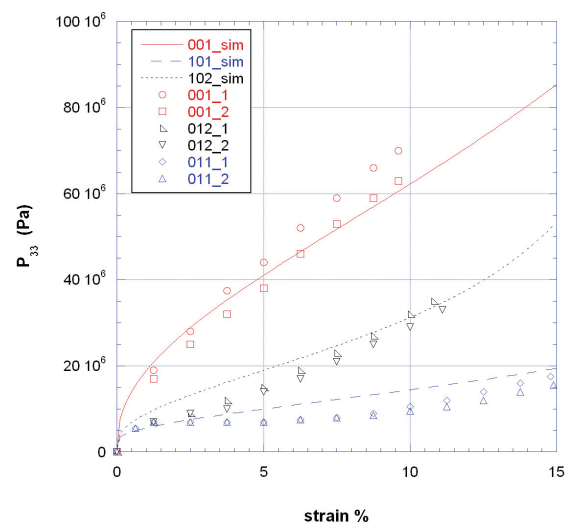
Funding Acknowledgments

LANL Directed Research and Development Program

Modeling Metallic Single Crystal Plastic Hardening through the Evolution of Dislocation Subgrain Structure

Benjamin L. Hansen, Curt A. Bronkhorst, T-3; M. Ortiz, California Institute of Technology

Fig. 1. shows a comparison between simulated and experimental stress-strain curves for copper single crystals.



A single crystal plasticity theory was formulated using sequential laminates to model subgrain dislocation structures. It is known that local models do not adequately account for latent hardening, as latent hardening is a nonlocal property as well as a material property (e.g., grain size and shape). The addition of the nonlocal energy from the formation of subgrain structure dislocation walls and the boundary layer misfits provide both latent and self-hardening of crystal slip. Latent hardening occurs as the formation of new dislocation walls limit motion of new mobile dislocations, thus hardening future slip systems. Self-hardening is accomplished by evolution of the subgrain structure length scale. The substructure length scale is computed by minimizing the nonlocal energy. The minimization of the nonlocal energy is a competition between the dislocation wall and boundary layer energy. The nonlocal terms are also directly minimized within

the subgrain model as they impact deformation response. The geometrical relationship between the dislocation walls and slip planes affecting dislocation mean that free path is accounted for, giving a first-order approximation to shape effects. A coplanar slip model is developed due to requirements when modeling the subgrain structure. This subgrain structure plasticity model is noteworthy as all material parameters are experimentally determined, rather than fit. The model also has an inherit path dependency due to the formation of the subgrain structures. Validation is accomplished by comparison with single crystal tension test results.

Evolution of the width between dislocation

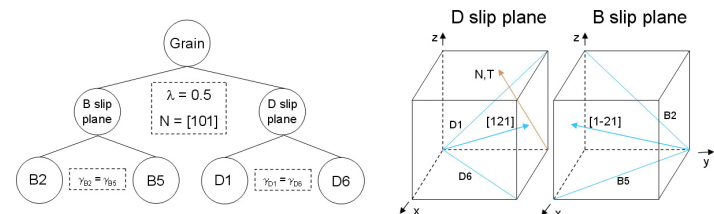


Fig. 2. The subgrain structure formed during the [101] simulation and the active slip systems in the crystallographic coordinates are given.

wall subgrain structures provides hardening mechanisms akin to those of self and latent hardening. The dimensions of the subgrain structure are determined by minimizing the nonlocal energy. The widths are equilibrated at each deformation step. Self-hardening occurs as the laminate width decreases due to increasing boundary layer energy of the laminate microstructure with increased deformation. The crystal bifurcating into separate regions to accommodate activation of new slip systems causes latent hardening to occur. The hardening behavior was validated against experimental tests of single crystal copper. Figure 1 shows a comparison between simulated and experimental stress-strain curves for copper single crystals. Three different crystal alignments are shown: [001], [101], and [102]. The crystal direction is aligned to the tensile axis. The single points are the experimental data from Franciosi (1985) [1]; the lines are computer simulations using the laminate hardening model. Note that there are two sets of experimental data for each crystal orientation. The simulations provide good validation for the hardening mechanisms of the model for these highly symmetric crystal orientations.

The subgrain structure formed during the [101] simulation and the active slip systems in the crystallographic coordinates are given in Fig. 2. The [101] orientation forms a simple laminate involving the B2 and B5 slip systems and the D1 and D6 slip systems.

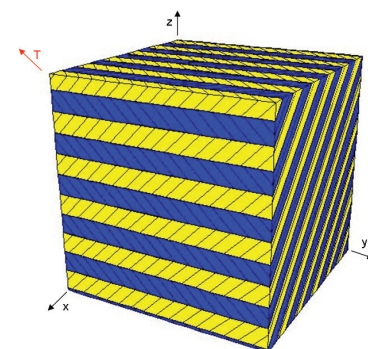


Fig. 3. Laminate formation

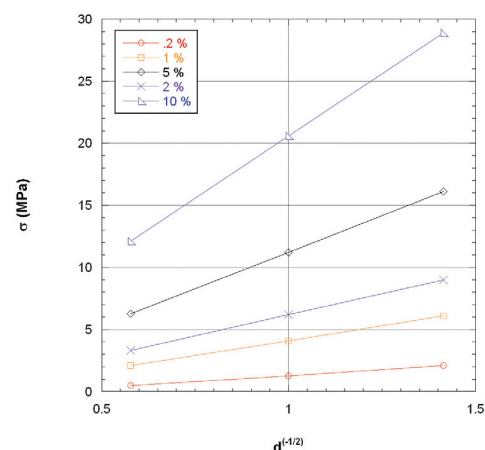


Fig. 4. Yield strengths

The black planes within the colored regions show the orientation of the active slip plane. The tensile axis, T, is included for reference.

All four active slip systems have the same strain level. The stress state of the sample has no shear stresses that would provide driving force for activation of further slip systems. The only nonzero stress in the simulation is the P33 component. Calculation of the Schmidt factor for the [101] orientation confirms that the four predicted slip systems have the highest Schmidt factor and should activate first. This is in agreement with Kalidindi and Amand [2], in which a traditional single crystal model was compared to a [101] copper single crystal deformed in compression. Kalidindi and Anand [2] also predicted that the same four slip systems would activate with equal strains. The deformation involves a shortening of the sample 2 direction (y in crystallographic coordinate system), while the 3 direction (or tensile axis, T, in the crystallographic coordinates) is elongated, and the 1 direction does not change. This same pattern of deformation is also seen in [2] with the experimental copper [101]-oriented compression sample. The trends are reversed due to the sample being in compression instead of tension (i.e., the 2 direction lengthens, while the 3 direction is compressed, and the 1 direction does not change in length). The agreement with deformed shape provides further evidence that the slip systems are correctly predicted. The stress response, activated

The dislocation wall normal, N, is parallel to the tensile axis, T – meaning the laminates form perpendicular to the tensile direction. A visualization of this structure in the crystal coordinates is given in Fig. 3 with color representing the regions in which the two different slip planes are active.

slip planes, and deformed shape are in excellent agreement with experimental data.

Figure 4 shows the yield strength of the [102] crystal response against the inverse of the square root of the initial grain size computed at various offset strains to the linear elastic deformation. The Hall-Petch relationship holds as the plots are linear. This relationship holds as long as a single layer of laminate forms, as was the case in all experiments examined here. It should also hold for higher-order laminate structures, except during the time steps in which the new laminates form. Uchic [3] and Greer [4] have recent work on the deformation response of micron-sized columns of single crystals. Their work showed that single crystals exhibit a Hall-Petch relationship, despite Hall-Petch having been found for polycrystals. The theory presented here shows the Hall-Petch effect due to the formation of dislocation subgrain structure with a grain size dependency, hence it provides a possible explanation for the Hall-Petch effect for both single crystals and polycrystals. Both Uchic[3] and Greer [4] also predicted a transition as sample size decreased to a breakaway flow behavior. This laminate theory would predict that transition as the length scale at which the deformation process no longer favors the formation of dislocation structures, but instead favors activation of a single slip system. The Uchic [3] experimental data is on a Ni₃Al-Ta alloy and the Greer [4] data is on gold, so no direct comparison is possible for the copper considered here. Thus, the theory presented here should follow the same trends discovered on micron-sized single crystals.

The results of simulations varying the grain shape are shown in Fig. 5. As this grain shape effect is due to the formation of subgrain structures inside the single crystal, rather than compatibility constraints within a polycrystalline sample, comparison would need to be made to similarly shaped single crystal tests. As of yet, no information is apparent for comparisons.

The model here contains only seven material parameters. All material parameters were measured directly by independent experimentation. It is

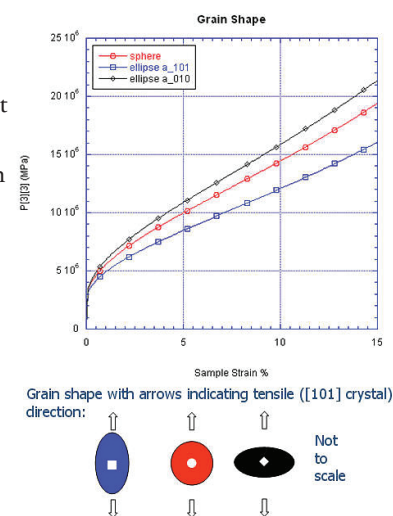


Fig. 5.

remarkable that the stress strain response can be predicted so well by the model with no fitting parameters.

For more information contact Benjamin L. Hansen at bhansen@lanl.gov.

- [1] P. Franciosi et al., *Acta Metall.* **33**, (9) 1601-1612 (1985).
- [2] S. Kalidindi, L. Anand, *Metall. Trans. A* **24A**, 989 (1993).
- [3] M. Uchic et al., *Scr. Mater.* **51**, 801 (2004).
- [4] J. Greer, W. Oliver, W. Nix, *Acta Mater.* **53**, 1821 (2005).

Funding Acknowledgments

DOE, NNSA Advanced Simulation and Computing (ASC) Program

Statistical Modeling for U-Nb Aging and Lifetime Prediction

Geralyn Hemphill, CCS-6; Robert Hackenberg, MST-6

Fig. 1. Universal age-hardening plot of U-5.6Nb TE data (solid symbols—fitting data from non-banded U-5.6Nb, open symbols—validation data from banded U-6Nb) and model fit to data (solid line). Alternate x-axes showing equivalent times at different temperatures are also shown. For reference, 10^5 minutes \approx 69 days and 10^7 minutes \approx 19 years.

The Enhanced Surveillance Campaign of NNSA's Advanced Simulation and Computing program supports many basic scientific studies that help stockpile stewardship. One important study concerns aging during long-term stockpile storage of the uranium alloy 6 wt% niobium (U-6Nb) [1,2]. Long-term aging may change the microstructure and properties of U-6Nb alloy components in ways adversely affecting performance. Traditional approaches to modeling based on fundamental physics are not feasible due to limitations in the understanding of the scientific fundamentals of age hardening [3,4]. Statistical modeling provides a way to develop

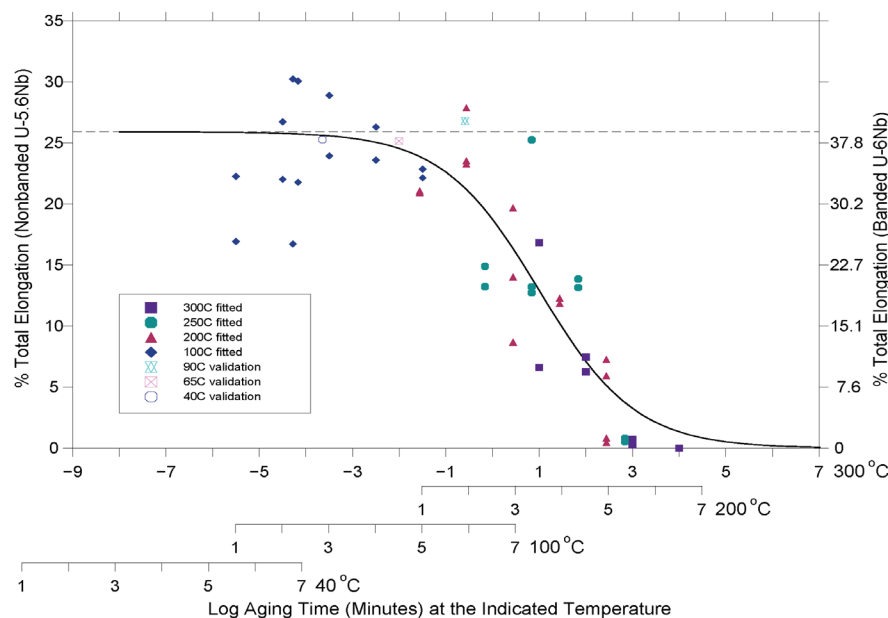
empirical predictive models for such data where first-principle models for complex physical mechanisms are not available.

We developed statistical models relating age-sensitive properties to time and temperature, drawing upon a large body of artificial aging data obtained from nonbanded U-5.6Nb and U-7.7Nb material [5]. These two nonbanded alloys represent the average (U-5.8Nb) and upper limit (U-8Nb) of compositions present on 100- to 200- μ m length scales in the banded U-6Nb material. There was concern that the high Nb bands would age faster and cross the failure threshold sooner than the mean or lower Nb bands. Thus, the high Nb (U-7.7Nb) band behavior would be the life-limiting factor for banded U-6Nb.

The aging response was tracked in U-5.6Nb and U-7.7Nb indirectly through the observed time-dependent changes in various mechanical properties. The following age-sensitive properties were measured at ambient temperature following artificial ages at 100°C, 200°C, 250°C, and 300°C: total plastic elongation, uniform elongation, first-yield strength, first-yield modulus, second-yield strength, ultimate tensile strength, and Vickers hardness.

Figure 1 shows our approach of collapsing the data from the four accelerated aging temperatures onto a universal aging response. This unique approach is based on the assumption of Arrhenius behavior governing the equivalencies of various time-temperature combinations. A separate term is included in the model to quantify the shift from one time-temperature domain to another. This allows the effects of both time and temperature on aging response to be assessed simultaneously.

The model also enables predicted curves to be produced at temperatures other than those at which data was obtained. Figure 2 provides an example of the model fit and extrapolations to lower temperatures for the property total elongation. In this plot, the time-temperature curves are plotted individually for each aging temperature. For each of the predicted curves in the plot, 95% confidence intervals are provided.



Reasonable model fits to artificial aging data for each of the properties in the U-5.6Nb and U-7.7Nb alloys were obtained, although the models fit some of the measured properties better than others. Useful age-sensitive property predictions were obtained for most of the elastic, strength, ductility, and hardness properties studied. The U-5.6Nb models were more robust and therefore are expected to have better predictive power than those of the U-7.7Nb models, especially at the lower aging temperatures of interest. Model extrapolations to longer times (up to 5 years) and lower temperatures (as low as 40°C) than those used for the model fitting agreed well with most of the validation data gathered for both nonbanded alloys, as well as banded U-6Nb, giving provisional validation of the fitted models.

In addition, the lifetime for each alloy at a variety of aging (storage) temperatures was evaluated, based on the failure criterion for the property total elongation. The error associated with the model fits provides 95% confidence intervals, which give the upper- and lower-bound lifetime estimates. From Fig. 2, the lifetimes can be obtained by finding the point at which a given predicted curve crosses the failure threshold (the yellow line).

Research such as this is essential for the understanding of our aging weapons systems. Studies concerning fundamental materials and material properties are crucial to decision making regarding the stockpile.

For more information contact GERALYN HEMPHILL at gsh@lanl.gov.

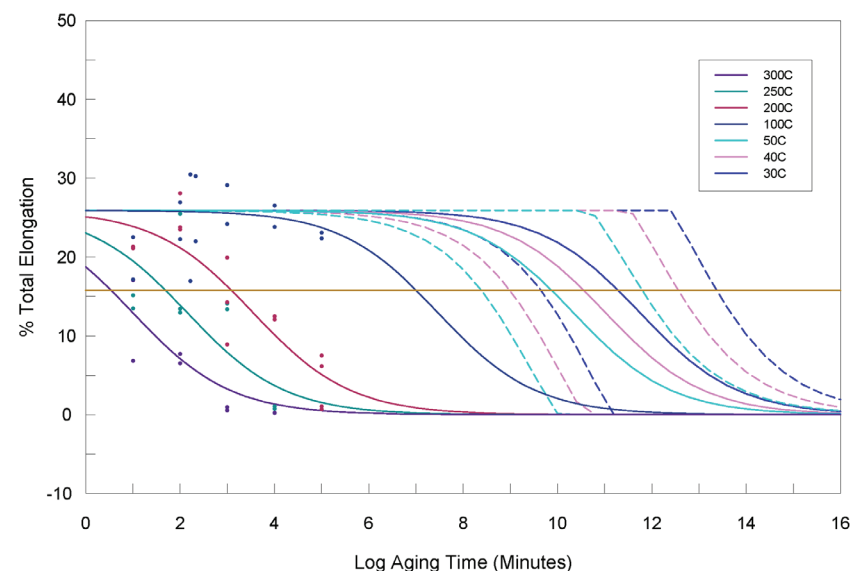


Fig. 2. U-5.6Nb TE fitting data (points), model fits to data at the artificial aging temperatures employed (solid lines), and low-temperature model predictions (solid line—mean, dashed lines—95% confidence intervals). The yellow line is the failure threshold.

- [1] R.E. Hackenberg, G.M. Hemphill, "Property and Lifetime Prediction in Aged U-Nb Alloys: A Statistical Assessment," LA-14389 (2009).
- [2] G.M. Hemphill, R.E. Hackenberg, "A Statistical Model for the Study of U-Nb Aging," *28th Compatibility, Aging and Stockpile Stewardship (CASS) Conf. Proc.* (2009) LA-UR-09-04733.
- [3] R.E. Hackenberg et al., *U-Nb Aging Final Report*, LA-14327 (2007).
- [4] A.J. Clarke, R.D. Field, R.E. Hackenberg, et al., "Low Temperature Age Hardening in U-13at.%Nb: An Assessment of Chemical Redistribution Mechanisms," *J. Nuclear Materials*, **393**, 282 (2009).
- [5] R.E. Hackenberg et al., "Synthesis and Characterization of Non-banded U-Nb Plate Material," LA-14316 (2007).

Funding Acknowledgments

- DOE, NNSA Advanced Simulation and Computing (ASC) Program
- DOE, NNSA Enhanced Surveillance Campaign

A Computational Study of the Ballistic Performance of Composite Materials

Balaji Jayaraman, Xia Ma, Paul T. Giguere, Duan Z. Zhang, T-3

In this study we assess the difference of ballistic projectile penetration between the traditional armor steel and a composite material. These simulations were carried out using CARTABLANCA [1], an object oriented, Java-based capability. CARTABLANCA is an Eulerian-Lagrangian code that uses the Material Point Method (MPM) along with a multipressure-based multiphase flow capability. Here, we use an Eulerian mesh with Lagrangian particles to handle multimaterial interactions as shown in Fig. 1.

Since this is an exploratory study, we assume that the projectile is a cylindrical-shaped tungsten bullet traveling at a speed of 2 km/s before impact on a stationary target. The projectile is 2 cm long and ~ 0.5 cm in diameter. To study the different deformation dynamics we consider two choices of the target. The first is a steel plate of 2-cm thickness and 5-cm long in each transverse direction. The second target is also chosen to be of the same dimensions, but is constructed as equal part of steel and composite layered on top of each other. In this comparison, we assume that the steel and composite are perfectly bonded. The dynamic interaction between the projectile and the target block is shown in Figs. 2 and 3 for the steel and steel-composite target respectively. In these figures the red particles represent the projectile (tungsten) while the blue (composite) and green (steel) particles represent the target block.

The constitutive relation for both the steel and tungsten projectile are described using the Johnson-Cook model [2], in which the yield stresses of the metals are functions of the effective plastic strain and the temperature. The model parameters are the same as used

by Zhang et al. [3]. For the composite we use an anisotropic linear elastic stress model along with the appropriate failure and damage criteria.

For more information contact Balaji Jayaraman at jbalaji@lanl.gov.

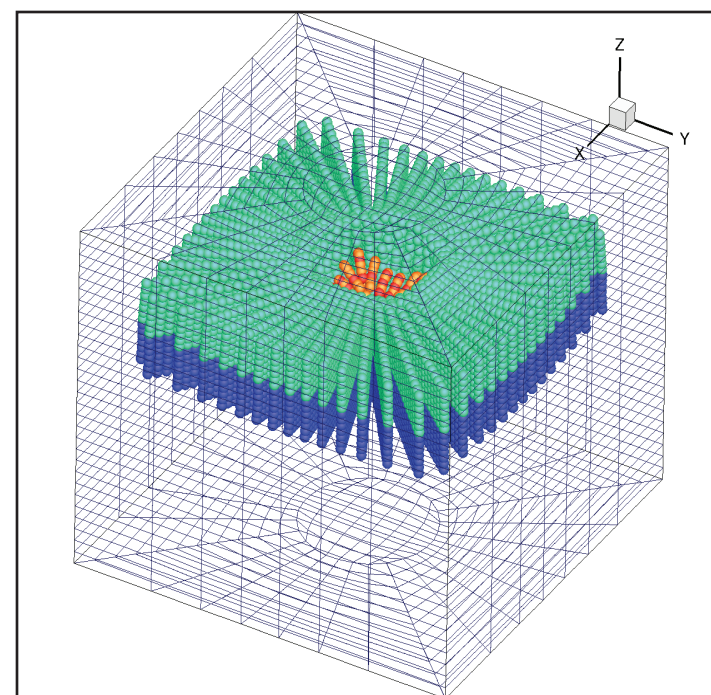


Fig. 1. Problem computational setup shows the background Eulerian mesh with the Lagrangian material points. Red particles depict the projectile while the light blue/green and dark blue particles represent the steel target and the composite material respectively.

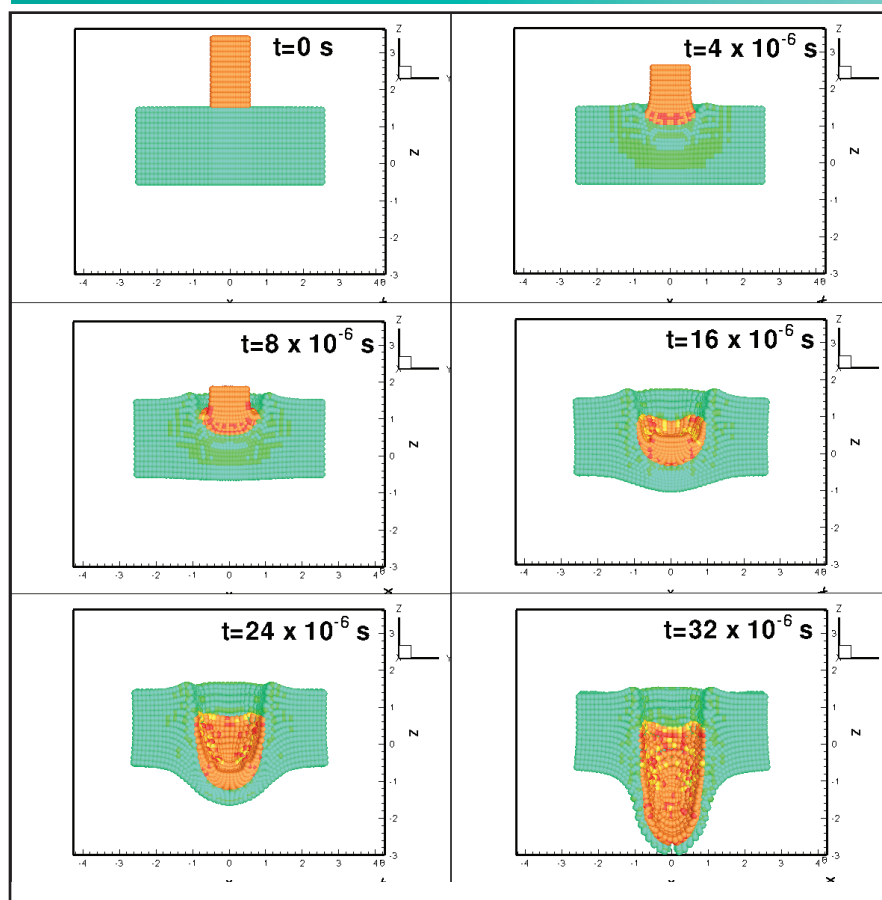


Fig. 2. Time evolution snapshots of a tungsten projectile penetration on a steel target block. The snapshots show a cross section at the center of the block.

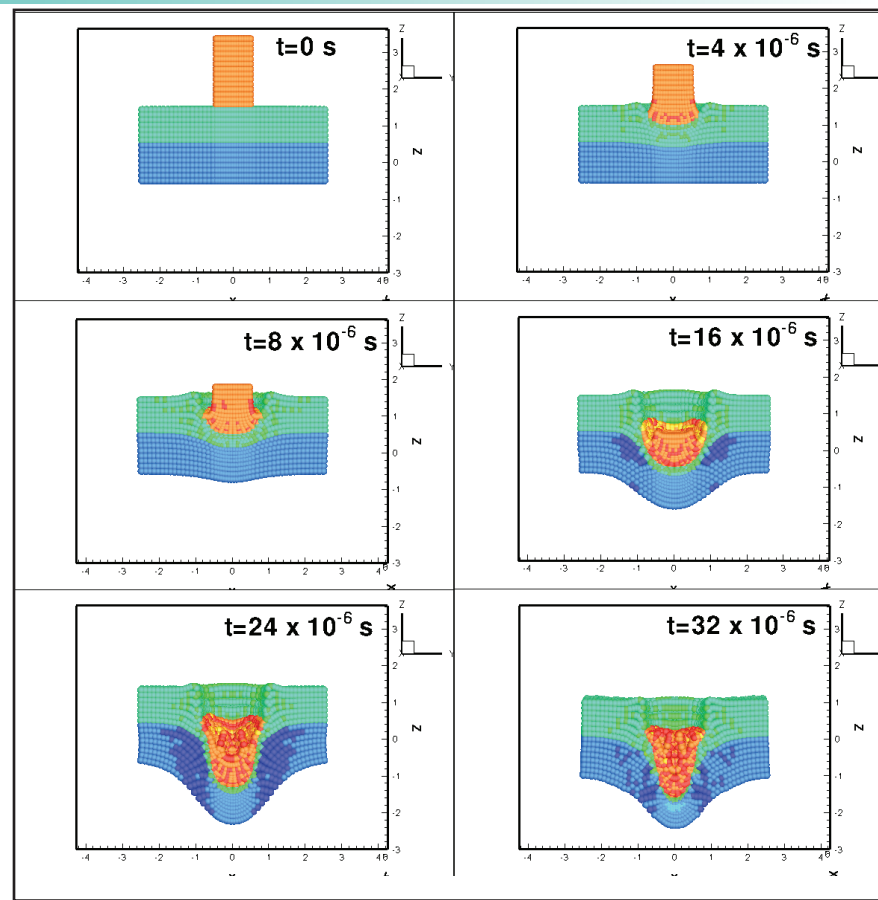


Fig. 3. Time evolution snapshots of a tungsten projectile penetration on a steel-composite material target block. The snapshots show a cross section at the center of the block.

- [1] CARTABLANCA, <http://www.lanl.gov/projects/CartaBlanca/>
- [2] G.R. Johnson, W.H. Cook, *Eng. Fract. Mech.* **21**, 31 (1985)
- [3] D.Z. Zhang et al., *J. Comp. Phys.* **227**, 3159 (2008).

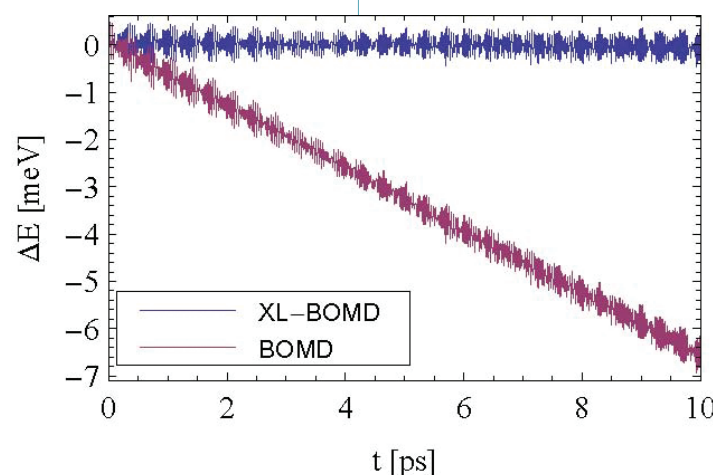
Funding Acknowledgments

DOE/DOD Joint Munition Program

Extended Lagrangian Born-Oppenheimer Molecular Dynamics

Anders M. N. Niklasson, T-1; Peter Steneteg, Linköping University, Sweden;
Marc Cawkwell, T-1

Fig. 1. The fluctuations in total Born-Oppenheimer energy ($\Delta E = E^{\text{tot}}(t) - E^{\text{tot}}(t_0)$) for a microcanonical eight-atom Si supercell simulation using XL-BOMD and regular BOMD ($\delta t = 1$ fs). The implementation was performed in the Vienna *ab-initio* simulation package (VASP), a plane wave pseudo-potential program for calculations based on Density Functional Theory (DFT).



Born-Oppenheimer molecular dynamics (BOMD) [1], based on the first principles of quantum mechanics, is currently the gold standard for a large class of atomistic simulations in materials science, chemistry, and molecular biology. Unfortunately, BOMD is often limited by some fundamental problems including a very high computational cost, unbalanced phase space trajectories with unphysical hysteresis effects, numerical instabilities, and a systematic long-term energy drift [2,3]. We have developed an extended Lagrangian BOMD (XL-BOMD) that overcomes some of these fundamental shortcomings [4,5].

The Born-Oppenheimer approximation includes a separation between the nuclear and the electronic degrees of freedom, where the forces acting on the atoms are calculated at the electronic ground state. The force evaluation is computationally very expensive, because of the demanding nonlinear self-consistent field (SCF) optimization of the electronic ground state. The number of SCF

iterations necessary to converge the solution dominates the computational cost and may range from ten to hundreds of cycles, depending on the system. To reduce the cost, the electronic solution is propagated between time steps. This means that some combination of the optimized electronic degrees of freedom, $\Psi^{\text{SCF}} = \{\psi_n^{\text{SCF}}\}$, from previous time steps is used as an initial guess in the SCF optimization,

$$\text{i.e.,} \quad \Psi^{\text{SCF}}(t) = \text{SCF} \left[\sum_{n=1}^N c_n \Psi^{\text{SCF}}(t - n\delta t) \right] \quad (1)$$

Propagation reduces the number of SCF iterations and thus the computational cost, typically by an order of magnitude, and is necessary for most practical simulations. Unfortunately, evolving the electronic solution through the nonlinear SCF procedure in equation (1) leads to an irreversible propagation, causing unphysical phase space trajectories with a systematic long-term energy drift [2,3]. Only by increasing the degree of SCF convergence, which increases the computational cost, can the systematic errors be reduced though the problems never fully disappear, since the SCF optimization in practice is always incomplete and approximate. Often the unphysical behavior of BOMD is ignored. By using thermostats, e.g., an artificial interaction with an external heat bath, the systematic errors in BOMD are covered up and may never be noticed. However, thermostats require a physically correct underlying dynamics and the errors are therefore never removed.

Our new XL-BOMD solves these fundamental problems by allowing a geometric integration of both the nuclear and the electronic degrees of freedom that preserves geometric properties of the exact flow of the underlying dynamics. This enables highly efficient and accurate long-term microcanonical simulations without a systematic long-term energy drift.

First principles BOMD based on density functional theory (DFT) [1] is given by the Lagrangian:

$$\mathcal{L}^{\text{BO}}(\mathbf{R}, \dot{\mathbf{R}}) = \frac{1}{2} \sum_i M_i \dot{R}_i^2 - U(\mathbf{R}; \rho^{\text{SCF}}) \quad (2)$$

where $\mathbf{R} = \{R_i\}$ are the nuclear coordinates and the dot denotes the time derivative. The potential $U(\mathbf{R}; \rho^{\text{SCF}})$ is the ground state energy including ion-ion repulsions calculated from DFT at the self-consistent electron density given by the wave functions, i.e., $\rho^{\text{SCF}} = \sum_n |\psi_n^{\text{SCF}}|^2$.

In XL-BOMD [4,5] we extend the dynamical variables of the Born-Oppenheimer Lagrangian with a set of auxiliary wave functions

$\Phi = \{\phi_n\}$ in harmonic oscillators centered around the evolving self-consistent ground state wave functions $\Psi^{\text{SCF}}(t)$

$$\begin{aligned} \mathcal{L}^{\text{XBO}}(\dot{\mathbf{R}}, \mathbf{R}, \Phi, \dot{\Phi}) = & \mathcal{L}^{\text{BO}} + \frac{1}{2} \sum_n \mu_n \int |\dot{\phi}_n|^2 d\mathbf{r} \\ & - \frac{1}{2} \sum_n \mu_n \omega_n^2 \int |\psi_n^{\text{SCF}} - \phi_n|^2 d\mathbf{r}. \end{aligned} \quad (3)$$

Here μ_n and ω_n are fictitious mass and frequency parameters for the harmonic oscillators. The Euler-Lagrange equations of motion in the limit $\mu_n \rightarrow 0$ are given by

$$M_k \ddot{R}_i = - \frac{\partial U(\mathbf{R}; \rho^{\text{SCF}})}{\partial R_i} \quad (4)$$

$$\ddot{\phi}_n(t) = \omega_n^2 (\psi_n^{\text{SCF}}(t) - \phi_n(t)) \quad (5)$$

The auxiliary wave functions $\Phi(t)$ are dynamical variables and can be integrated by a geometric integration scheme, for example, the time-reversible Verlet algorithm. $\Phi(t)$ will stay close to $\Psi^{\text{SCF}}(t)$, since the auxiliary wave-functions evolve in harmonic wells centered around the ground state solutions. Using $\Phi(t)$ as an initial guess in the SCF optimization, therefore provides a very efficient SCF procedure within a framework that can preserve geometric properties of the flow of the underlying dynamics. This is in contrast to the conventional propagation shown in equation (1), where the nonlinear and approximate SCF procedure breaks the time reversibility. Simulations based on XL-BOMD can incorporate time reversibility and other exact properties of the underlying dynamics, even when the force evaluations are based on an incomplete and approximate SCF optimization. This greatly enhances the efficiency and accuracy of MD simulations.

The XL-BOMD approach is quite general and can be used in combination with higher-order symplectic integration schemes, for the propagation of wave functions (as presented here), density matrices [4], effective single-particle Hamiltonians, and the density [6]. In contrast to the popular Car-Parrinello molecular dynamics scheme [1,7], the atomic trajectories evolve on the Born-Oppenheimer potential energy surface with the total self-consistent

Born-Oppenheimer energy (kinetic + potential) as the constant of motion. This gives a more accurate dynamics and allows for longer time steps in the integration. Figure 1 illustrates the total energy conservation of XL-BOMD in comparison with regular BOMD with the SCF optimization initialized with a higher-order wave function interpolation from previous time steps. The SCF convergence criteria was the same for the two examples, but the computational cost is lower for XL-BOMD and the systematic energy drift is removed. Figure 2 shows another example implemented in LATTE, a self-consistent tight-binding program developed at LANL [6]. Long-term energy stability is achieved even using only one SCF cycle per time step.

In summary, XL-BOMD provides a new generation of high-performance BOMD within a rigorous and transparent theoretical framework that can combine accuracy and long-term stability with a low computational cost.

For more information contact

Anders M. N. Niklasson at amn@lanl.gov.

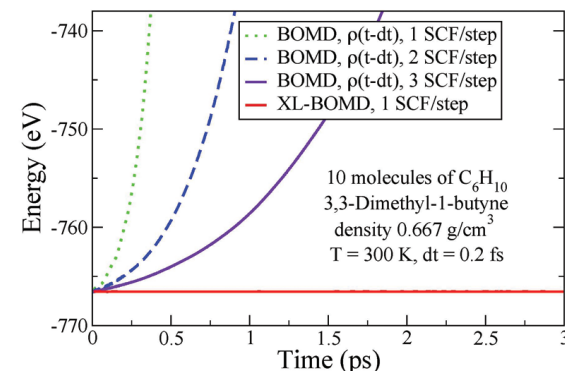


Fig. 2. The total Born-Oppenheimer energy for XL-BOMD and regular BOMD based on a simple propagation of the electronic charge, $\rho(t - \delta t)$, from the previous time step. The implementation was performed in LATTE [6].

- [1] D. Marx, J. Hutter, *Modern Methods and Algorithms of Quantum Chemistry*, Institute for Computing (2000).
- [2] P. Pulay, G. Fogarasi, *Chem. Phys. Lett.* **386**, 272 (2004).
- [3] A.M.N. Niklasson, C.J. Tymczak, M. Challacombe, *Phys. Rev. Lett.* **97**, 123001 (2006).
- [4] A.M.N. Niklasson, *Phys. Rev. Lett.* **100**, 123004 (2008).
- [5] A.M.N. Niklasson et al., *J. Chem. Phys.* **130**, 214109 (2009).
- [6] E.J. Sanville et al., *LATTE: Los Alamos Transferable Tight-binding for Energetics*, LA-CC 10004 (2010)
- [7] R. Car, M. Parrinello, *Phys. Rev. Lett.* **55**, 2471 (1985).

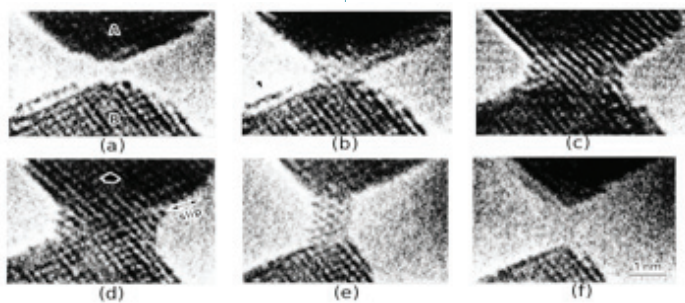
Funding Acknowledgments

LANL Directed Research and Development Program–Exploratory Research (ER)

Simulating the Mechanical Behavior of Metallic Nanowires over Experimentally Accessible Timescales on Roadrunner

Danny Perez, T-1; Chun-Wei Pao, Academia Sinica, Taiwan; Sriram Swaminarayan, CCS-2; Arthur F. Voter, T-1

Fig. 1. A series of high resolution transmission electron microscope (HRTEM) images showing the contact formation, retraction, and rupture processes of two gold tips. (a)-(c): contact formation process with (a) corresponding to $t=0$ s, (b) to $t=110/30$ s, (c) to $t=129/30$ s; (d)-(f) retraction and rupture processes: (d) $t=0$ s, (e) $t=2$ s, (f) $t=3$ s. Figures from [1].



For many years now, we have witnessed explosive growth in our capability to control the structure of materials down to the atomic scale. For example, it is now possible to bring the tip of an atomic force microscope, which might be only a few atoms wide at the apex, in contact with a surface. In the case of a metal tip contacting a metal surface or another tip, a bonded contact forms, and, if the tip is then lifted away from the surface, surface reorganization and diffusion often creates a nanowire that maintains a connection between the tip and surface. This process, as captured by high-resolution transmission electron microscopy, is shown in Fig. 1. This kind of manipulation can be exploited for the intentional creation and study of nanowires, whose width can sometimes be reduced all the way down to a single atomic chain by continuing the retraction process. These nanowires are an ideal probe of the nanoscale behavior of materials, be it mechanical [1] or electrical [2], and hence are of interest both for fundamental studies and because of their expected importance in various nanotechnology applications such as electrical conductors and electrical or mechanical switches. However, a deeper understanding of the fundamental nanoscale behavior

of materials is required before these applications can become widespread. Indeed, experiments at the nanoscale are hard to control and sometimes lack the resolution necessary to fully understand how the systems behave. For example, metallic nanowires are often seen to completely disappear from one

frame to another, leaving one completely in the dark about the basic mechanisms leading to their failure. Further, the act of imaging itself leads to an uncontrolled increase of the temperature of the wire, making it very difficult to perform these experiments in controlled and reproducible conditions.

There is thus a pressing need for atomistic numerical simulations to complement such experiments and help interpret and understand them. The most powerful tool for performing this kind of simulation is molecular dynamics (MD), whereby one integrates the equations of motion of all the atoms in the system, advancing the positions and velocities of the atoms by repeatedly taking small steps forward in time. In this way, one learns about the evolution of the system with full atomistic detail. However, for many systems and processes we would like to study there is a serious problem with the mismatch in time scale. For example, the nanowire-stretching process discussed above usually takes place over seconds or, at the very fastest, milliseconds. In contrast, conventional MD simulations are limited to a time scale of about one microsecond, even on the fastest parallel computers, i.e., 10^3 to 10^7 times faster than the experimental reality. Because of this extremely large time gap, physical arguments suggest that current simulations do not adequately represent reality.

The way to overcome these limitations is to use so-called Accelerated Molecular Dynamics (AMD) methods to reformulate the problem in a form that is more amenable to computer simulation. For example, the Parallel Replica Dynamics (ParRep) [3] method developed at LANL generates a proper evolution of the system while allowing a parallelization of the problem in the time domain. This allows one to make optimal use of massively parallel computers to reach time scales that are orders of magnitude longer than what could be done with conventional MD. When implemented on petascale supercomputers like Roadrunner, the ParRep method enables one to study the evolution of nanoscale systems (containing about a thousand atoms or so) at the unprecedented rate of about 0.1 ms per wall-clock hour when using 12,000 replicas on 120,000 Opteron and IBM Cell Broadband Engine (Cell BE) cores, thereby allowing a direct connection between experiments and fully atomistic simulations.

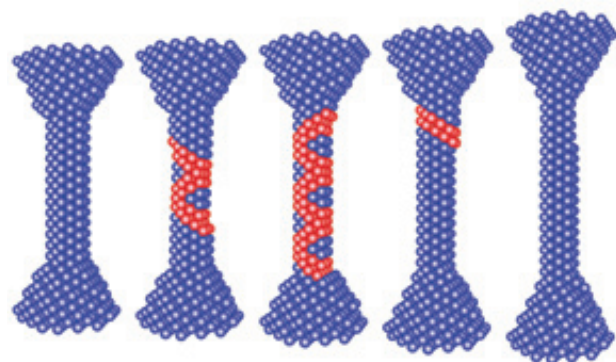


Fig. 2. Early stage of a ParRep simulation of the stretching of a silver nanowire on Roadrunner at a temperature of 300K and a retraction velocity of 10^{-5} m/s. From left to right, $t=0\ \mu\text{s}$, $t=30\ \mu\text{s}$, $t=60\ \mu\text{s}$, $t=90\ \mu\text{s}$, $t=150\ \mu\text{s}$. Atoms in noncrystalline configurations are shown in red.

Using ParRep on Roadrunner, we simulated silver nanowire stretching experiments similar to those illustrated in Fig. 1 for different nanowire sizes, temperatures, and retraction velocity, reaching more than one millisecond of simulation time in a few instances, more than a thousand times longer than conventional techniques would have allowed. Overall, we were able to study the change of behavior of these wires while varying the strain rate by more than four orders of magnitude, a feat that was unthinkable before the advent of Roadrunner.

Thanks to these simulations, a picture of the evolution of these systems is emerging. The basic plastic reaction of the system when subjected to strain is to create stacking faults along (111) planes. These stacking faults are highlighted in red in Fig. 2. The formation of a zig-zag network of such stacking faults causes the release of internal stresses while leading to the elongation and narrowing of the wire. Interestingly, almost all wires, independently of temperature or strain rate, initially behave this way. As the wire is stretched further, given enough time, these stacking faults annihilate, leaving behind a defect-free wire that is uniformly thinned down relative to the initial configuration. These simulations illustrate the unique ability of these nanostructures to, under suitable conditions, heal themselves when

subjected to severe external constraints. Note that this self-healing behavior emerges only on long time scales that are completely inaccessible to standard MD simulations.

The later stages of the simulation of the evolution of nanowires also revealed some completely unintuitive mechanisms by which plastic deformation occurs at the nanoscale. As shown in Fig. 3, one of these mechanisms is the conversion of bulk-like segments of the wire into low-symmetry helical structures (here a fivefold-symmetric icosahedral structure). These structures appear to be extremely tolerant to mechanical constraints. Indeed, the conversion process between two relatively stable conformations offers a continuous pathway for the stretching to occur, in contrast with competing mechanisms that lead to the accumulation of defects and ultimately to failure. Through this pathway, elongations in excess of 100% have been observed without failure. Once again this process, involving structures that are not allowed by the symmetry of the bulk crystal, demonstrates the unique ability of nanoscale systems to react to their environment in completely unintuitive ways.

The conjunction of innovative algorithms and methods with the unprecedented computational power of Roadrunner enables us to simulate, for the first time, the mechanical behavior of metallic nanowires, which are widely foreseen as playing a major role in the next generation of nano-devices, on experimentally accessible time scales. With this new capability it is now possible to directly assist in the interpretation of experiments as well as in the design of novel structures with precisely tailored properties.

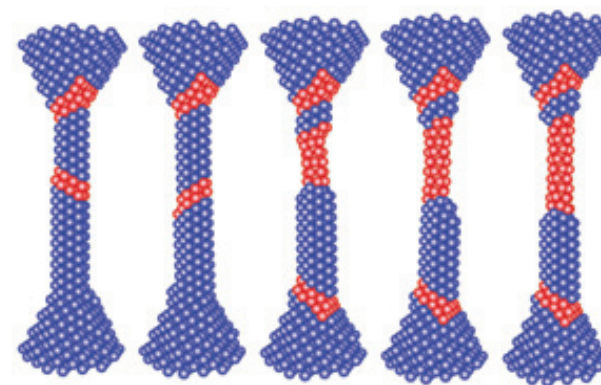


Fig. 3. Late stage of a ParRep simulation of the stretching of a silver nanowire on Roadrunner at a temperature of 300K and a retraction velocity of 10^{-5} m/s. From left to right, $t=165\ \mu\text{s}$, $t=180\ \mu\text{s}$, $t=195\ \mu\text{s}$, $t=210\ \mu\text{s}$, $t=225\ \mu\text{s}$. Atoms in noncrystalline configurations are shown in red.

For more information contact
Arthur F. Voter at
afv@lanl.gov.

- [1] W. Liang, M. Zhou, F. Ke, *Nano Letters* **5**, 2039 (2005).
- [2] J.I. Pascual et al., *Science* **267**, 1793 (1995).
- [3] T. Kizuka, *Phys. Rev. Lett.* **81**, 4448 (1998).
- [4] A.F. Voter, *Phys. Rev. B* **57**, 13985 (1998).

Funding Acknowledgments

- LANL Directed Research and Development Program
- DOE Office of Basic Energy Sciences (BES)

Author Cross References

A

Ahrens, James.....	48
Albright, Brian J.....	10, 36
Anderson-Cook, Christine	108
Andrews, Malcolm J.....	26
Ankeny, Lee.....	48
Aronson, Igor.....	16

B

Bakosi, Jozsef.....	2
Barber, John L.	68
Batista, Enrique R.	66, 114
Batista, Victor S.	114
Bent, John.....	96
Bergen, Ben.....	10, 36
Berlyand, Leonid.....	16
Berman, Gennady P.....	4
Bhattacharya, Suman.....	48
Bhattacharya, Tanmoy	76
Bingert, John F.	116
Bishop, Alan R.	4
Bock, Nicolas	70
Bowers, Kevin J.....	10, 36

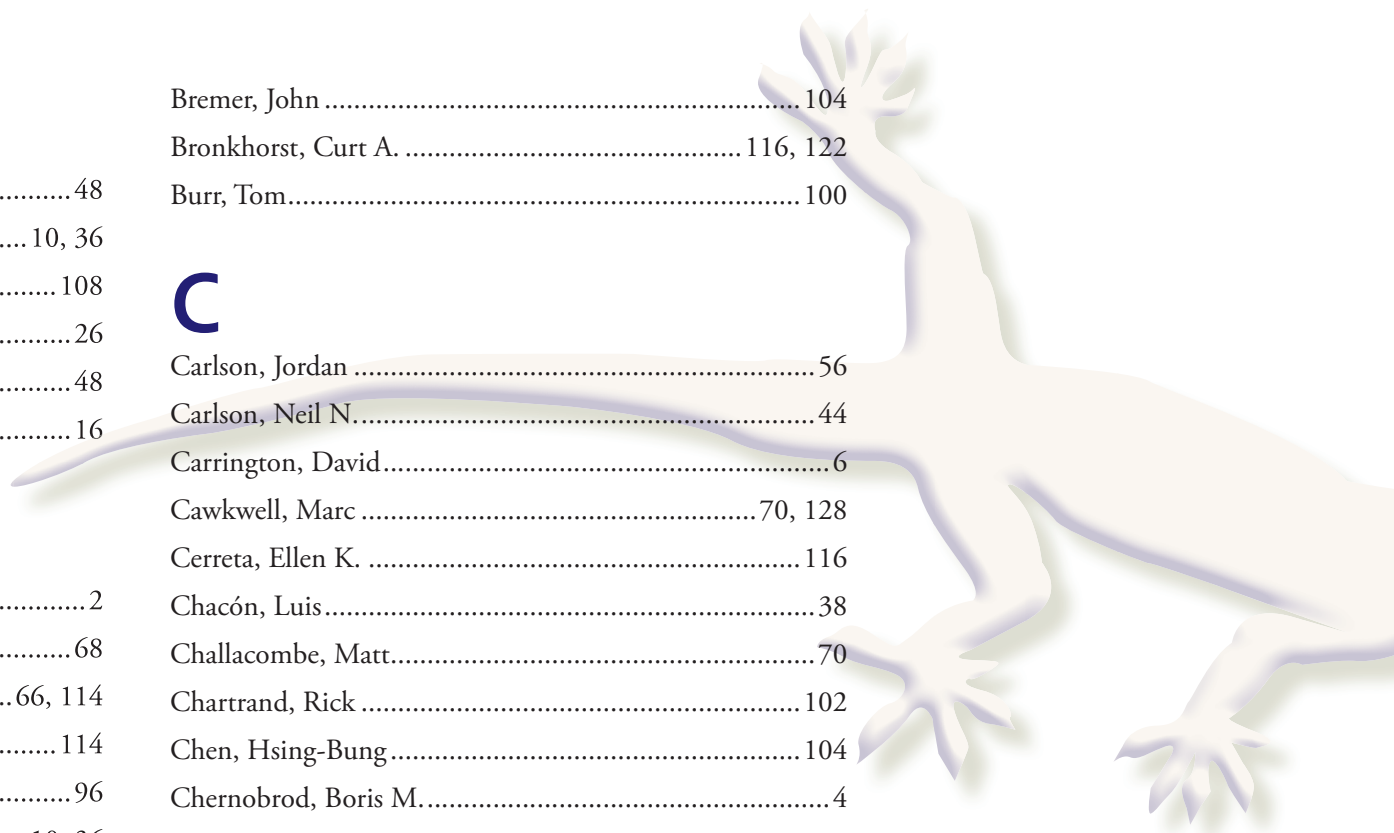
Bremer, John	104
Bronkhorst, Curt A.	116, 122
Burr, Tom.....	100

C

Carlson, Jordan	56
Carlson, Neil N.....	44
Carrington, David.....	6
Cawkwell, Marc	70, 128
Cerreta, Ellen K.	116
Chacón, Luis.....	38
Challacombe, Matt.....	70
Chartrand, Rick	102
Chen, Hsing-Bung.....	104
Chernobrod, Boris M.....	4
Chumak, O.O.....	4
Clements, Bradford E.....	118
Colgan, James	72

D

Dai, William W.	8
Daligault, Jerome	54
Daniel, David.....	48, 56
Daniels, Marcus	76



Daughton, William S.	10
Dean, Sumner	24
DeBardleben, Nathan	106
Delzanno, Gian Luca	38
Dendy, Joel.....	42
Densmore, Jeff	60
Desai, Nehal.....	48
DeVolder, Barbara	60

E

Eidenbenz, Stephan.....	110
-------------------------	-----

F

Fasel, Patricia.....	48, 56, 92
Fichtl, Chris A.....	40
Finn, John M.	38, 40
Francois, Marianne M.	12, 44
French, Alfred D.	78
Fryer, Chris L.	52, 60
Fugate, Michael.....	108

G

Garimella, Rao, V.	14, 28
Germann, Timothy C.	120
Gibson, Garth	98
Giguere, Paul T.....	126

Glimm, James	30
Gnanakaran, S.....	76, 78
Gorshkov, Vyacheslav N.	4
Graves, Todd	108
Greene, Bob	52
Grider, Gary.....	98, 104
Gupta, Sanjib	54
Gyrya, Vitaliy.....	16

H

Habib, Salman	48, 50, 56
Hackenberg, Robert	124
Hamada, Michael.....	100, 108
Hansen, Benjamin L.	116, 122
Heitmann, Katrin.....	48, 50, 56
Hemphill, Geralyn	108, 124
Hengartner, Nicolas	100
Hennelly, Scott P.	80
Henning, Paul J.....	60
Higdon, David	50, 108
Hsu, Chung-Hsing	48
Hungerford, Aimee	52, 60
Huzurbazar, Aparna V.	108

J

Jakubikova, Elena	66, 114
Jayaraman, Balaji	126
Johnson, Glenn	78

K

Kelley, Timothy M.	20, 60
Kelly, Elizabeth	108
Klamann, Richard	108
Korber, Bette	76
Kucharik, M.	28

L

Langan, Paul	78
Lawrence, Earl	50
Li, Hui	58, 62
Li, Shengtai	32, 58
Lipnikov, Konstantin	16, 34
Lipscomb, William H.	88
Livescu, Daniel	18, 20, 24
Lowrie, Robert B.	22, 90
Lukić, Zarija	48, 56
Luo, Sheng-Nian	120

M

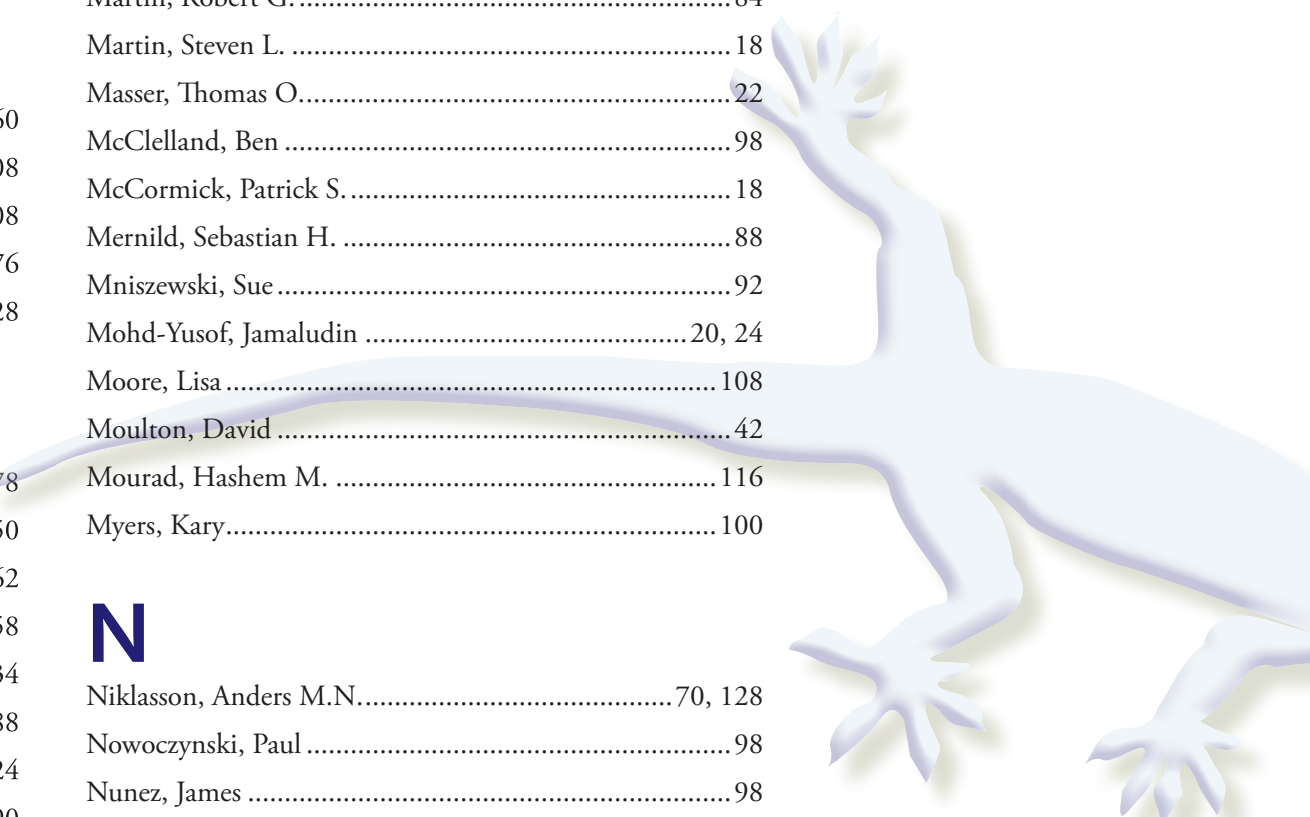
Ma, Xia	126
Mark, Graham	48
Martin, Richard, L.	66, 114
Martin, Robert G.	84
Martin, Steven L.	18
Masser, Thomas O.	22
McClelland, Ben	98
McCormick, Patrick S.	18
Mernild, Sebastian H.	88
Mniszewski, Sue	92
Mohd-Yusof, Jamaludin	20, 24
Moore, Lisa	108
Moulton, David	42
Mourad, Hashem M.	116
Myers, Kary	100

N

Niklasson, Anders M.N.	70, 128
Nowoczynski, Paul	98
Nunez, James	98

O

Ortiz, M.	122
----------------	-----



P

Pao, Chun-Wei.....	130
Perez, Danny.....	66, 130
Petersen, Mark R.....	18, 24
Picard, Richard.....	100, 108
Polte, Milo	98
Pope, Adrian	48, 56

R

Raskin, Cody	52
Ringler, Todd D.	90, 94
Ristorcelli, J.R.....	2
Rockefeller, Gabriel	52, 60
Rollin, Bertrand	26
Rosner, Judah L.....	84
Ross, Amy R.....	116
Roytershteyn, Vadim	10

S

Sanbonmatsu, Karissa Y.....	80
Sanchez, Kathryn	104
Sanville, Edward.....	70
Schofield, Samuel P.....	28
Scott, Cody	104
Sharp, David H.	30
Shashkov, Mikhail J.....	28

Shen, Tongye.....	78
Snoeberger III, Robert C.....	114
Springer, Everett.....	92
Steneteg, Peter.....	128
Svyatskiy, Daniil.....	34
Swaminarayan, Sriram.....	120, 130
Swartz, Blair K.	12

T

Ticknor, Lawrence.....	108
Timmes, Frank.....	52
Torous, S.V.....	4
Torres, Aaron	104
Torres, David	6
Tung, Chang-Shung.....	82
Turley, Milton	104

U

Urbatsch, Todd.....	60
---------------------	----

V

Vander Wiel, Scott	108
Vassilevki, Yuri	34
Vivoni, Enrique.....	92
Voter, Arthur F.	66, 130

W

Wagner, Christian	50
Wall, Michael E.....	84
Wendelberger, Joanne.....	108
White, Amanda.....	92
White, Martin.....	50, 56
Williams, Brian	50, 108
Wingate, Meghan.....	98
Wohlbier, John G.	22

X

Xu, Hao	62
---------------	----

Y

Yan, Guanhua	110
Yin, Lin.....	10, 36

Z

Zhang, Duan Z.	126
---------------------	-----

